

From: Tom Eich [tbe@microunity.com]
Sent: Sunday, April 30, 1995 2:55 PM
To: tbr (Tim B. Robinson)
Subject: Re: Pandora module mechanical design meeting

>Tom I just found this. My absense (assuming you went ahead) should not
>be taken to indicate a lack of interest! Please let me know if there
>are any issues remaining that I can help with.

>
>Sorry to have dropped the ball on this.

>
>Tim

>
>Tom Eich wrote (on Mon Apr 24):

>
> Hi,

>
> I would like to meet with you to review mechanical details of the common
> feature set to be used in the Hermes modules. These details have an
> impact on the pcb layouts of the four module types (Euterpe, Mnemo,
> Calliop, PCI/Hermes bridge), and also on the backplane. The Euterpe and
> Mnemo layouts have been completed to the level where the pcb routing
> constraints are well understood, and so we can now evaluate the mechanical
> design options that will accomodate these layout constraints.

>
> Can we meet on Wednesday, at 3:00pm in the boxers conference room? Because
> I would like to do some real-time concept selection (from design approaches
> I'll have available), I would like to limit the attendance, but if you want
> anyone else to be there, please let me know.

>
> -Tom

> Tom Eich	tbe@microunity.com
> MicroUnity Systems Engineering, Inc.	
> 255 Caspian Dr. Sunnyvale, CA 94089	
> (408) 734-8100, (408) 734-8136 fax	

No problem, David, Philip, and I went ahead with the concept review and selection. I sent out a brief summary subsequently and we'll review the selected design tomorrow.

-Tom

Tom Eich	tbe@microunity.com
MicroUnity Systems Engineering, Inc.	
255 Caspian Dr. Sunnyvale, CA 94089	
(408) 734-8100, (408) 734-8136 fax	

From: tbr
Sent: Sunday, April 30, 1995 2:53 PM
To: Tom Eich
Subject: Pandora module mechanical design meeting

Tom I just found this. My absense (assuming you went ahead) should not be taken to indicate a lack of interest! Please let me know if there are any issues remaining that I can help with.

Sorry to have dropped the ball on this.

Tim

Tom Eich wrote (on Mon Apr 24):

Hi,

I would like to meet with you to review mechanical details of the common feature set to be used in the Hermes modules. These details have an impact on the pcb layouts of the four module types (Euterpe, Mnemo, Calliop, PCI/Hermes bridge), and also on the backplane. The Euterpe and Mnemo layouts have been completed to the level where the pcb routing constraints are well understood, and so we can now evaluate the mechanical design options that will accomodate these layout constraints.

Can we meet on Wednesday, at 3:00pm in the boxers conference room? Because I would like to do some real-time concept selection (from design approaches I'll have available), I would like to limit the attendance, but if you want anyone else to be there, please let me know.

-Tom

Tom Eich
MicroUnity Systems Engineering, Inc.
255 Caspian Dr. Sunnyvale, CA 94089
(408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: tbr (Tim B. Robinson)
Sent: Sunday, April 30, 1995 11:00 AM
To: paulp (Paul Poenisch)
Cc: geert
Subject: processing of Castor/Pollux and Calliope1

Paul Poenisch wrote (on Fri Apr 21):

Calliope-1

The data used to generate this reticle set is not quite as upto date as Orchis (which is also not up to current standards) but it should be usable. Unfortunately when we processed the first lots of this device through the middle metal layers (metal 2 through 3) we discovered that the reticle vendor that made this set failed to hold the dimensional tolerance needed for this process. As a result we are having problems processing this device.

We believe that we can use the current reticle set to produce some Calliope-1's but the photomasking engineers will have to hand carry the lots through the metal layers. This will result in some slowing of the lots when compared to Orchis (which does not have this problem) but should still allow us to get enough devices out to allow debug of the device.

Paul, this seems a little different from what was said in the meeting last week. If the problem is just poor processing of the masks, then why would we expect that necessarily to carry over to mnemo an euterpe. (You may remember the issue was Al saying he would not process these if they had the same problems.) I realize there are other reasons we have to change, but I would like to be sure I fully understand the reasoning.

Tim

From: wingard (Drew Wingard)
Sent: Saturday, April 29, 1995 3:43 PM
To: microlunatics
Subject: SITN Seminars for the week of May 1

Here are this week's EE310 (Advanced Interconnect Systems) and EE380 (Advanced Parser Generators) Seminar notices.

Bill Z. reports that the EE310 seminar did not show up tape delayed on Tuesday evening as published in the Spring Quarter SITN guide.
Does anyone have the updated broadcast schedule?

Best Regards,
Drew

EE 310 Seminar
Tuesday, May 2, 4:15 pm, McCullough 134

Materials and Integration Issues in Advanced Interconnect Systems

John Sanchez
Paul Besser

Advanced Process Development, Integrated Technology Division Advanced Micro Devices,
Sunnyvale, CA

ABSTRACT

The development of interconnect materials systems for advanced VLSI circuit technologies poses significant challenges. The choice of materials and interconnect architecture which satisfy device performance requirements is relatively straightforward. However it is much more difficult to choose the materials and processing schemes which allow for efficient manufacturing and which also provide high yields and sufficient device reliability. Often these materials and processing choices involve metallurgical issues related to the choice of Al alloy, deposition conditions, and the barrier or anti-reflection layers. These metallurgical issues often affect other concerns such as RIE etch patterning and line resistance. We provide several examples where the control of processing variables and the choice of interconnect materials can be utilized to optimize manufacturing, yield and interconnect reliability. We describe the methodology for controlling CuAl₂ second phase formation during Al-Cu alloy sputter deposition for the purpose of optimizing the performance during RIE etch patterning. Next, the effects of alloying additions to Al (such as Cu and Si) on Ti + Al => TiAl₃ reaction kinetics is described. We provide a method for utilizing the beneficial aspects of Ti underlayers on Al interconnects while minimizing the Al consumption by TiAl₃ formation. Finally, we characterize resistance changes during accelerated electromigration testing, and describe the effects of metallurgical evolution (such as second phase precipitation) on interconnect resistivity. The effects of barrier/shunt layers on resistance increases as the Al interconnect voids is also DESCRIBED.

BIOGRAPHY:

John Sanchez, Jr., is a Senior Member of the Technical Staff in the Advanced Process Development-Films Group, Integrated Technology Division, at Advanced Micro Devices, Sunnyvale, CA. His primary responsibilities include the development of metallization systems for future (0.25 μ m and beyond) logic and memory technologies, with particular emphasis for "designing in" manufacturability, process yield and reliability. Other activities include thin film reactions, metallurgical evolution in thin films, mechanisms of electromigration and stress-induced voiding, and the effects of film microstructure on stresses. His educational background includes a B.S., M.S. and Ph.D. in Metallurgy/Materials Science at University of California, Berkeley. Previous positions include AMD (Supervising Engineer), Xerox Palo Alto Research Center (Member Research Staff), and Max-Planck Institut for Metal Research in Stuttgart, Germany (Visiting Scientist).

EE380 Computer Systems Colloquium

Spring Quarter 1994-1995

Lecture #5

Date: Wednesday, May 3, 1994

Time: 4:15-5:30 pm

Location: Terman Auditorium

SITN: Thursday, Channel E3, 8:00-9:15

Speaker: Terence John Parr

Title: This is not your father's parser generator

Abstract

LL(k) and LR(k) parsers were rigorously defined in the 1960's, however, LL(1) was the only practical parsing strategy as even LR(1) was intractably large. Finally, LALR(k) was defined around 1970 and the world was happy--LALR(1) was practical and much stronger than LL(1). A proliferation of LALR(1) tools followed with YACC and clones dominating parser generation completely--parsing theory has barely budged ever since.

YACC is tired. No good object-oriented interface has been defined and programmers complain about it's lack of flexibility and features. Further, LALR parsers are difficult to debug, hard to fully understand, and are sensitive to action placement (i.e., you can introduce an ambiguity into the grammar with an action). Most importantly, today's truly nasty parsing problems, such as C++, cannot be adequately solved using 20 year old technology. A leap in technology analogous to the jump from LL(1) to LALR(1) is required.

This talk introduces ANTLR (a component of PCCTS), a public-domain parser generator that combines the flexibility of hand-coded parsing with the convenience of a parser generator. ANTLR provides "predicates" which let the programmer systematically direct the parse via arbitrary expressions using semantic and syntactic context. ANTLR also integrates the description of lexical and syntactic analysis, accepts LL(k) grammars for k>1 with extended BNF notation, provides sophisticated parser error handling, and can automatically generate abstract syntax trees.

Biography

Terence Parr is president of Parr Research Corporation, a software development and consulting company located in Minneapolis, MN, and is a part-time postdoctoral fellow at the Army High Performance Computing Research Center in Minneapolis. He is the primary author of the public domain PCCTS language tool kit, which is used extensively throughout the research and industrial community. A book about PCCTS is in progress.

Terence received a BS in Computer Science in 1987 and a PhD in Electrical Engineering in 1993 both from Purdue University. His thesis involved the efficient construction of exponentially-complex parsers; most of his friends are still wondering how he got a thesis on parsing theory past an electrical engineering thesis committee.

Contact Information:

Terence John Parr
parrt@parr-research.com

From: graham (Graham Y. Mostyn)
Sent: Friday, April 28, 1995 6:58 PM
To: tbr
Cc: graham
Subject: Need to strategize mask changes

Tim, I believe we need to hold a strategy meeting to discuss the selection and timing of mask revisions for Calliope0, Calliope1 and Castor/Pollux. (An important part of the equation is that the CAD portion of these activities should be coordinated with - and not adversely impact - CAD activities for the Euterpe/Mnemo tapeout. While circuit design work could start now, I see a bottleneck at the CAD stage later.)

Before calling such a meeting, however, I'd like to hear your thoughts on the following. Perhaps this session should be followed with a broader-scope discussion on scheduling and prioritization of all MUSE mask set tapeouts.

Objective: Define and schedule Pollux and Castor mask changes.

Invitees: paulp, mudge, hopper, bill, geert, + staff

- Select DRCs for correction on Pollux (metal and perhaps diffusion).
- Provision of SOFA ring oscillator on Pollux, requested by fab.
- Addition of new metal waffle rules to Pollux and Castor to allow manufacture.
- Consider necessary changes to Castor
- Schedule of circuit design work, CAD activities and tapeout; relation to Euterpe/Mnemo mask scheduling.

Graham.

From: hopper (Mark Hofmann)
Sent: Friday, April 28, 1995 1:54 PM
To: hardheads
Subject: Allegro licenses

Hi,

We have just 3 Allegro licenses. With the board revision and tapeout activity, as well as cross-probing between PCB and schematic, competition for licenses is high. We are looking into the purchase of more copies of the software. In the meantime, this is a plea: Please if you are not actively using the tool could you exit (rather than just iconifying, or leaving it running) so others can get a chance to use the tool?

-thanks,
hopper

From: craig
Sent: Friday, April 28, 1995 1:00 PM
To: craig@mnemosyne
Subject: Reminder- tapeout review

** Calendar Appointment **

Date: 4/28/95
Start: 11:00 am
End: 12:00 pm
What: tapeout review

From: tbr (Tim B. Robinson)
Sent: Wednesday, April 26, 1995 10:43 PM
To: euterpe
Subject: Known problem list

We are in the final stages of tapeout at last. We have a list of known open issues, most very minor or obscure, which we do not plan to correct before tapeout. I would like to review this list with those interested on Friday at 11am, Engineering conference room, to determine the impact of these anomalies, if any last minute corrective action is necessary, and how best to document the remainder.

The current list is attached below:

V Need verification to run with multicycle RAM models with accurate X injection on read-write collisions and badly shaped write cycles.

S uu/uumemuv has merged its hiccup & machine check versions of store-load conflict by assuming that ooo cases trying to cause str-ld or other hiccups will instead machine check. But no such machine check logic exists, allowing strange behavior, although avoidable, to occur without direct warning. In other words, we depend on the programmer to avoid certain "abnormal" DBuffer/Dtag read/write collisions, as discussed below the threshold of consciousness. We should either take a machine check if they happen or prevent them from trying to hiccup, because ooo preempts have no PC to which to hiccup.

S We should add a cerberus bit to disable branch prediction in case there is a bug we need to disable in real hardware.

H We decided sometime back to reduce permutation uncertainty on signals propagating through synchronizers by sharing 1 cmos2ecl synchronizer per cmos signal. Are there any left?

S We have to invent algorithms for changing unsynchronized CMOS controls to ECL logic or we can gate them with something like mem mngmnt enbl. Much of cerberus octlet 6 needs this. For example, there are not even synchronizers on NB's lplevel, and we probably don't want' to force with mem mngmnt enbl because it makes testing the lplevel restricted to one mode. The GA field can be safely changed when mem mngmnt enbl is off. The cache size fields must only be changed when NB is empty and if all (data and instruction) memory accesses have LVA<47> set until the update must have safely happened. This means code and data must come from buffer. Note that it is impossible to read back the change in the usual way during this interval because the cerberus address does not have bit 47 set, so a delay loop will have to be used. The NB priority (number of entries for low priority) field must only be changed when no uncached-offchip nor cache-miss ifetches are possibly being processed. This must be enforced across all cylinders until a readback confirms cerberus has the new setting. The easiest way is to run lBuffer.

S Gating ICache size with mem mngmnt enbl (suggested above) does not create the following case but makes it more likely (however, the delay loop solution above should also cover this case): If a jump fetches its target in lFe/ICC with iOffChip==1 but then, after the ISRAM is read with an original cache size (perhaps forced), mem mngmnt is enabled in time for atprchk to send cacheableOrNoAllocR11 to ICC, then iccxci17 will think it is OK to use wrongly indexed ISRAM data as a target ICache hit. Data will be wrongly indexed if the cache size changed along with the rest of the cerberus octlet 6 write to enable mem mngmnt if the

cache size changed from any unforcing when mem mgmnt mode enables.

H Silly Logic: Do we have a check for hrbuf's in tau distribution? Mws found a plain ff in rgxmit (actually two in series, a waste). If topt is timing two ticks between hr's without checking clock distribution, then we need to run the gloss check. Rgxmit & HZ are now fixed, others remain?

S To avoid a deadlock between a cyl holding the cache controller but also wanting the xcptn reg, and a cyl oppositely wanting and holding the same resources, we claim "forward progress" and release CC when an instruction does not complete but instead hiccups because the exception register is busy.

The scenario is:

Cylinder 0	Cylinder 1
-----	-----
Cause exception.	ICache miss on ifetch.
Enter event mode.	Grab CC; fill; no "forward progress"
Handler gets IorD cache miss.	yet, so we still own CC.
Hiccup while waiting for CC.	1st instruction just filled gets xcptn.
Xcptn reg now deadlocked.	Hiccup waiting for xcptn reg.
	CC now deadlocked.

We must not claim "forward progress" on general hiccups because the cache miss recovery causes hiccups and we want the hiccuping instruction to retry before its cache lines could be stolen.

We break this rule to add locked-out exceptions to the cases claiming "forward progress"; this increases the risk of thrashing leading to a livelock hang, but that may not be too likely because the job waiting to use the xcptn reg has proven itself to not need CC a 2nd time to finish anyway. Thus only an ICache miss is possible after the hiccup, and the cylinder locking the xcptn reg can probably progress through the thrashing to eventually release the xcptn reg.

S Maybe we are supposed to check Cache Coherence Required (an exception) on the (I or D) cache line to be replaced and if that tag has its CC bit set, raise the corresponding exception? We currently only report Cache Coherence Required on cache hits to the non-replacement access, which is of limited usefulness. We currently report the same 0,1,2,3 access types for read,write,execute,gateway as we do for other memory exceptions instead of the special CC Required 0,1,2,3 access types for read,write,replacement,undefined.
We may have a similar problem on the FVA, because a Cache Coherence Required may expect to report the address from the tag to be replaced instead of the address of the miss.

S A software/architectural consensus is emerging that GTLB detail exceptions (as well as CC required?) should have been higher priority than illegal physical address. It seems current priority was an accident.

S We should add a cerberus bit to disable machine checks?

V From jeffm Fri Apr 7 14:57:22 1995

Subject: Re: Test status - exlltest3

Mark Semmelmeier writes:

> > From jeffm Fri Apr 7 14:28:12 1995

> > When trying to access data cached, when the dcache size is zero, and

> > the dtags are X, the test goes to X.

> Do we have to make this work? We have all these places where

> cache size is (or is going to be) gated by memory management

> mode and other places (like checking tags) gated by memory

> management mode and we may have trouble gating

> with the OR of the two without gating races? Also if we

> cant figure out a way to gate once at the front of the

> piped copies, we have to hit more blocks with fixit logic.

> Maybe its not a big deal, but I hope we can keep this low

> priority.

I would agree - as long as it works with tag hit or miss (i.e. for real). I can certainly update the test to check both cases.

S Store&Swaps (synch ops) integrity may be damaged by backdoor DTag accesses. I think I can say that the fill case should be impossible because CC would have to have been busy and locking the cache line, thus preventing the earlier swap busy injection (awkward proof). See uuprblmr8.Veqn and swapstore section of notes in uustepuu.pla.

H We need to verify that paths that are declared DC but which are really just multi-cycle are not too slow for the cycles allocated for transition.

H Silly Logic: RGPC sends PL to its muxenff (2 gate levels) for linkage read but also sends out of data path to RGXMIT making TOpt think the muxen is at the end of the longer run. Maybe a buffer to isolate RGXMIT is appropriate.

H Silly Logic: UU sends UUvldGoUV straight out of a pla's flop for long run to ES, but then sends same through buf+orff, which forces pla to power up enough to have driven buf+orff as if over in ES. Probably a 2nd flopped copy of vldGoUV would be cheaper.

- - t h r e s h h o l d - - o f - - c o n c i o u s n e s s - - - - -

M If the SOFA clock to DRAM clock ratio is set at less than 10:1 and if at the same time the SOFA clock to Hermes clock ratio is 1:1 (actually 2:1 since the Hermes clock is double edge active), then the combined bandwidth of a series of back to back read responses on both Hermes channels simultaneously, at the same time as DRAM reads are returning will overrun the prb. Something will get dropped and there is nothing in the design to catch the failure.

Explanation:

The prb (peripheral return bus) has a maximum data bandwidth of 8 Bytes / 4 ticks = 2 B/t. (B is Byte, t is tick, the sofa clock period).

The maximum (load) bandwidth from all the non-slow peripherals (HC0, HC1, and DRAM) must be ≤ 2 B/t.

Eq. 1) (max load BW HC0) + (max load BW HC1) + (max load BW DRAM) ≤ 2 B/t.

When HC0, HC1 and DRAM are all returning data on the prb at the same time, the prb is time-division-multiplexed with fixed slot allocations. These slots are allocated in the ratio 2:2:1 to HC0:HC1:DRAM. So...

Eq. 2a) (max load BW HC0) $\leq 2/5 * 2$ B/t

Eq. 2b) (max load BW HC1) $\leq 2/5 * 2$ B/t

Eq. 2c) (max load BW DRAM) $\leq 1/5 * 2$ B/t

Now, if R is the ratio of DRAM clock to sofa clock, Rt is one DRAM clock period. For DRAM configurations which use all 32 data pins the maximum DRAM load bandwidth is 4 B/Rt. So using Eq 2c), $4B/Rt \leq 1/5 * 2$ B/t implies $R \geq 10$.

Note that if HC1 is not being used, its bandwidth is available for the DRAM to use.

S Sandeep and others want an ERes in event mode to just reenter event mode and not cause a machine check so they can use ERes as a breakpoint to debug event handler code.

S The TSA demands that GTLB and cache tag detail exceptions will be ignored on a the post-event-handler replay of the instruction causing the exception. This is complicated by possible hiccups on retry and the fact that the event handler may choose to dispatch to a new task which would _not_ want its first instruction ignoring detail exceptions.

S Event entry assumes no DBuffer read/write collisions during the save & restore sequence. Do not use another cylinder to store to a different cylinder's new PC/R1 hexlet if that latter cylinder could leave non-event mode to enter event mode. Should this be a machine check? There is no restriction for a cylinder storing to its own new PC/R1 hexlet.

- S UU/CC cannot recover from writeback reads that get write-read collisions in the D data SRAM. The collision is of hexlet resolution. The workaround is to not mess around with DBuffer stores in the region reserved to implement the Dcache (CC will not cause this case on its own). Should this be a machine check?
- S UU/CC cannot recover from snake D tag reads that get write-read collisions. The workaround is to be sure the cache controller and all other cylinders are not writing that line's tag in the vicinity of the D tag read. Should this be a machine check? This seems more likely to be one to let happen but ignore the error, possibly by looping with voting.
- S SMUX now produces the reserved instruction exception. We had a bug in the current implementation that SMux64 behaves the same as SMAS64; i.e. SMux64 wrote a destination register when it was not supposed to. Since the user can always restore SMAS64's data register to get the same effect as SMux64, SMux seems only an infrequently usable minor performance enhancement. Craig rebuts that: SMux is also useful for writing a subset of the bits of an octlet, as may occur in modifying a bit field in a structure; SMux also permits a much faster implementation like stores since there is no dest reg.
- S It may be possible to for some cylinders to gang up on another cylinder and cause so much store activity that the latter cylinder can make no progress due to store-load collision.
- S It may be possible to lock out interrupts from a cylinder by arranging interrupt-repelling multi-job (microcode/step) instructions to always be busy on eta3 (1 of 4 major cycles) when interrupts are checked (eta 3 since event entry 1st job must be eta3).
- S The Execution Privelege Level Required (XA) bits in the ITag are always filled with lowest privelege on a miss instead of the GTLB's field. Since IFetch still checks privelege vs. the GTLB on every branch taken and every sequential crossing of an I-architectural-page-size boundry, a programmer would have to set privelege changes in the GTLB finer than the I-architectural-page-size or do backdoor ITag stores to notice.
- S If the program jumps from (ibuffer/cachedOrNoalloc) onchip to (uncached) offchip but ICache is filled with an image of the offchip region (either via backdoor ITag or earlier execution with a different GTLB setup), then the first octlet of instructions at the target may be fetched from ICache instead of off chip. This can be considered forced-at-target no-allocate behavior. We might be able to save a few atoms in iccxc17.Veqn if this and other onchip-offchip transitions cause a hiccup.
- V Does exrlharder test vary the distance from R1 NB creation to event entry use?
- S Performance: Each ICache fill writes the ICache array 8 times for 8 octlets. Each of these writes potentially causes 3 IFetch reads from unpredictable cylinders to be stalled. There is a comparator on cache index[14:12] validating collisions so that at least IBuffer code IFetching will not be impaired by activity in uncontrolled other cylinders executing from ICache.
- S Performance:
If an instruction that has been fetched after an Icache miss has some kind of trouble in issuing or completing successfully, then there will be a performance hit beyond the best case which is the delay from the ICache access to the DCache access (about 5 major cycles).
Issue hold delays would be waiting for a specific slot (usually 1 but 3 for some instructions like synch ops) or waiting for a dependent src register. If NB is causing the register dependency, then the issue hold could be indefinite.

Instruction completion may be impeded due to Dbuffer/Dcache write-read conflicts, gtlb conflicts, and destination register anti-dependency with NB. The NB case can again be indefinite. I forgot to mention that some instructions, although not being impeded, may take a long time to complete just because they are microcoded. The longest of these is 64 bit multiply which occupies 23 issue slots before completing. We could do a special optimization that would recognize this as a no-D-memory instruction and release the ICache-locked-up-cache-controller sooner, but I am not sure it is worth the extra work and bug risk.

- S Performance: NB delays the release of the NB entry after its returned data has been read by the pipe if the data might need to be immediately written back into NB (e.g. for moving data to IFetch/ICache). Rather than use an instruction decode to decide this, NB uses the slot type such that any pipe read of NB in a store slot is assumed to need a delayed release. Thus there will be unnecessary release delays for loads that happen to occupy store slots, making NB appear fuller.
- S Performance: Cache controller could use other cylinder's load slots to make fill requests now that it has better handshaking with issue that lets it know when slots are available.
- S Performance: Since the cache controller only makes fill requests on otherwise unused load-type pipe slots, other background activity (preempts or cylinder restart after an interrupt) could significantly impede fill requesting, which affects the latency of all cylinders that use the cache controller.
- H Silly Logic: CC sends CCmissAcceptIR13 (ccstart imissout) to ICC which carefully tries to keep it halfswing. But CC also sends it to ccseq which is going to make it full swing!.
- H Silly Logic: CDIO write data in should be done with 1/4 rate hr's controlled by eta 0 not tau instead of half rate with front end 2muxes. The only reason not to is if TOpt would limit 1/2 to 1/4 paths to only 1 tick but we needed 2 ticks.
- H Silly Logic: The OR of vldSN128WrtIR11 and vldSN128WriQIR11 should be precomputed and sent into atpdcld which has 11 loads doing the same thing.
- H SR is used to copy AUndx1500R2 for IFe and thus keep 1 more load off that critical bus. This was done before HZ existed when SR was the best candidate (candidates had to already be needing the bus for some important reason). HZ is now available and is better placed to do it.
- H Silly Logic: Both HZ and AT are computing store write indexes. AT sends ATndx1406R12 to CC & LTstrWrNdxR12R13 to SR. HZ sends R7R8 to CC {va to match cc fzn_va} and HZndx1404R13R16 to CDIO and CTIOD. Looks wasteful, especially of wiring.
- H Redundant Logic: ES & UU overlap in validating conditional branches to branch. There may not be many gates because of the AND-OR structure in ES already needed to choose which conditional case naturally defaults to zero on non-branches. Does it select 1 on unconditionals?
- H Redundant Logic: CTIOD uses ~adrs13 (~ndx9) to suppress writing the dirty bit, but the AT store dirty write enable is already fully accurate. needs to be reevaluated and maybe ripped out.
- H Redundant Logic: Cache size decoding of the force into AUndx1500R2[14] is useless since all cache sizes < 32K force 14.
- S Performance: We could try to make Cache miss hiccuping more complicated so that the hiccup retry lines up with the resolution of the cache miss

to reduce the miss latency. But we would like to keep interruptible cache miss capability too.

S Performance: CC could augment its vamatch comparator circuitry and write the DTag at the same time as the 1st hexlet so that the miss data would be available piecemeal instead of only after all 4 hexlets are written. This helps miss latency in most cases. Hit would have to be restricted to the already filled hexlets.

S A special memory-map address to clear the exception lock for when an errant cylinder fails to do so would be nice. But gmo says as long as another cyl can tell who is at fault and still take interrupts, this would just be a luxury.

H Top-Level Pruned Logic: AUndx1500R2[1:0].

From: doi (Derek Iverson)
Sent: Wednesday, April 26, 1995 8:24 PM
To: guarino; gmo; gregg; lisa; jeffm; lisar
Cc: hestia
Subject: Software Bringup Meeting Minutes - April 26, 1995

Software Bringup Meeting

April 26, 1995

Next Meeting: May 3 at 2:00 pm. <-- Note Time Change!!!!

Attendees: guarino, gmo, doi, gregg, lisa, jeffm

New Action Items

Item: Modify startup code in stb/stand to read the cerberus node number
Who: gmo
Status: New

Review of Action Items

Item: Plan for testing remote debugging environment
Who: everyone
Status: Pending

04/13 This will be discussed at the next meeting.
04/20 Short discussion on the pieces required and if we want to have
a standalone version of the hostio software.
Wally has started on this already.
04/26 Any work with snoopy/dram models will begin after euterpe tapeout
when jeff has more time.

Item: Schedule a meeting to go over how to run stuff on HW simulator and
read HW traces.
Who: jeffm, lisar, doi
Status: Pending

04/05 doi is going to talk to lisar about when this would be appropriate.
lisar has volunteered to show up at the next meeting and discuss
this.
04/13 lisar presented a strategy that will give other the ability to
run and do software debug with the hardware simulators. Jeffm
to write up a crib sheet that will aid with debugging in
the hardware accelerator environment.
04/20 Dave Tomlinson will require such a session. Sometime after May 8.
04/26 No change.

Item: ukernel needs to detect machine checks and `do the right thing'
Who: gmo
Status: Pending

04/13 No new progress.
04/20 On list of things to do. Lower priority.
04/26 Ditto.

Item: Modify tests in diag tree to use tcc instead of tgcc
Who: guarino, jeffm, doi
Status: Pending

- 03/29 Loretta has changed the diag tests to use tcc instead of tgcc but has not checked in the changes. Lisar wants these changes to be made after we are able to run the tests successfully using tgcc.
- 04/20 Loretta has the stand/diag tests converted to use TCC. Derek is going to modify the Makefiles in stand/diag/memhi so that the programs that required tgcc will have an explicit rule.
- 04/26 Loretta is ready to check in the TCC related changes. Derek is to check with lisar if this is OK.

Suspended Items

Item: Unsnap code
Who: sandeep, guarino
Status: Suspended.

- 02/15 The issue of restarting the hardware from an IKOS dump was discussed and the need for an architectural snap/unsnap facility was questioned. Since the meeting it has been re-discovered (jeffm wasn't there to remind us of an earlier decision) that we are planning on loading architectural state into an IKOS simulation and not from a total IKOS logic dump. We also determined that when it came time to run some of the larger tests (real-time benchmark) we would need the capability to start an IKOS simulation from an architectural dump anyhow.
- 03/01 For the short term we are going to focus on a simpler approach for loading and running DVTs, the kernel, and kernel tests. This item will likely come back in April.

Item: Create performance test plan
Who: jeffm, guarino
Status: [11/30] No progress as focus is on functionality.

We continue to run tests to help us compare terp vs hardware performance.
We still need to put together the actual performance tests that need to be run on the hardware.

Completed Items

Item: Need test to demonstrate the 'no forward progress' condition in terp
Who: jeffm
Status: Done.

- 04/05 Jeffm is going to write 'cachenasty4' in the hopes that it will create the situation in which the current terp will hang. This will enable lisa to see if her 'fix' does. The sequence of events required to create the problem are:
1. icache miss in cylinder 'a'
 2. the next instruction gets a dcache miss in cylinder 'a'
 3. some other cylinder has an icache miss (alias) on the same line.
- 04/13 Jeff has written cachesyncnasty. There are still problems running

the test on terp, but not the right one.
04/26 The test caused terp to hang. Lisa has a solution.

Item: Terp Feature Status
Who: gmo
Status: Removed. New section started for this.

Item: Tests need to be written to verify performance issues
Who: lisar, claseman
Status: Removed. Added new section to track progress.

02/22 We need to flag performance problems as errors.
Tests could be identified (and perhaps written) to measure
and verify performance of the hardware for things like cache
misses, tlb initialization, exceptions, etc.
03/01 Lisar has started writing these tests.
03/08 Work continues.
03/15 Tim Claseman is assisting.
03/22 We need to generate a list of tests that we think should be written
first. Jeff suggested dcache fills, icache fills, dram and hermes
accesses.
03/29 Tim has come a long way up the learning curve and is now focused
on producing the first 4 tests that were requested.
04/05 Tim has queued a hermes access performance test on the zycad.
04/13 Time has run a test on the zycad but apparently wasn't configuring
the dram properly. This is now fixed and ready to run again.
04/20 Tim to check in the tests and the verification group will
run them along with all the other tests.
04/26 Still looking for the first test to be checked in.

Terp Feature Status

new o Add support for host I/O through the sdram
done? o Holes in address spaces => machine checks
inprog o Reflect "forward progress" change in the hardware
- believed done.
inprog o Ifetch protection granularity
- performance vrs accuracy tradeoff
inprog o Fetch instructions as octlets
inprog o Simulate Ifetch queue
o Accuray wrt HW simulator(s?)
o Better latency model for Calliope accesses
o Implement hardware configuration through Cerberus regs
(SDRAM paramters, dram.enable?)
o Checkpoints/Snapshots
punt o Model PCI
change o hermes and cerberus timeout machine checks
- does jeff have a test for this? Currently the TSA is
followed (immediate mchk) instead of waiting for watchdog
timeout.
change o ability for terp to load hermes sections
- lisar would like this functionality added

Performance Test Status

Tim should be ready to start checking in tests.

Test Status and General Discussion

The nullTest still hasn't run successfully on the IKOS. Problem related to the GVA changes.

The ltlbtran test goes to X.
The exception_prio_test has passed.

Question for Lisar:

When do we have a full calliope simulation available (IKOS)? This topic was raised as we talked about when the Snap/Unsnap item should be brought back from the Suspended list.

From: paulp (Paul Poenisch)
Sent: Wednesday, April 26, 1995 6:06 PM
To: tbr (Tim B. Robinson); geert (Geert Rosseel); bill (William Herndon); wingard (Drew Wingard); mudge (john mudge); herman; warren (Mark Warren); al (Albert Matthews)
Cc: paulp (Paul Poenisch); trancy (Trancy Tsao)
Subject: TAB bonding on Cronus

Hi,

Trancy and I have been discussing the results of the meeting with MMS yesterday. If it turns out that we indeed need on substrate capacitors for Cronus the size of the substrate will reach a size that may make it difficult to do the ILB TAB bonding on our current equipment.

We have done some rough estimates of the size of the substrate needed for various options:

Case	Die size	
	1.7 cm die	2.0 cm die
No capacitors, no under die coat:	1.8 cm substrate	2.1 cm substrate
No capacitors, under die coat:	2.0 cm	2.3 cm
1 capacitor, under die coat:	2.2 cm	2.5 cm
2 caps, opposit corn., under die coat:	2.4 cm	2.7 cm
2 caps, adjacnet corn., under die coat:	2.2 x 2.4 cm	2.5 x 2.7 cm
4 caps, under die coat:	2.4 cm	2.7 cm

For the case of no capacitors, no under coat and a 1.7 cm die we could still use a 48 mm TAB frame, for all other cases a 70 mm TAB frame will be needed. When the substrate starts to approach 1" (~2.3 cm) two more problems will show up for ILB. First the pedestal that holds the substrate will reach a size that will require a significant change in the way it is heated (redesign of that section of the system). Second the bond arm on our ILB TAB system will have a reach problem which will require a redesign of the TAB frame clamp mechanism, this problem may set an absolute limit on the size of the substrate (a number to be determined).

The timing of the cronus parts coming in for ILB TAB bonding is likely to conflict with the bonding of Euterpe and Mmemo. Because there are several changes to the ILB system that would be needed to bond Cronus we may not be able to do both types of chips in the same time period. Therefore we believe that it may be time to look into having the ILB work done by an outside vendor, preferably MMS.

We need to discuss this issue in the near future. If there will be a meeting to discuss other issues for Cronus packaging in the next few days we would like to have this added to the list of subjects. If there will be no meeting scheduled by the middle of next week we think we need one.

Paul, Trancy.

From: paulp (Paul Poenisch)
Sent: Wednesday, April 26, 1995 3:35 PM
To: hardheads
Subject: MobIMOS power routing

Today in a meeting on design rules and fracture flow we identified a serious power routing problem that effects all devices designed for MobIMOS to date (Calliope-0 through Euterpe and Mnemo). The problem is reliability related, all current designs can be expected to fail within a week, possibly as little as an hour, due to Vdd to Vss shorts caused by corrosion.

Background:

Early on it was decided to connect the seal ring of the die and space transformer to Vdd. This was based on a proposed layout of the space transformer that is now obsolete. Later when the first die was being layed out it was determined that a structure to prevent cracking of the dielectric layers at the edge of the die during wafer saw was needed (virtually all devices have such a structure). Because of the geometry of the situation this crack prevention structure must be connected to ground. At about the same time it was found that we had a power distribution problem for test and the seal ring was incorporated into the power distribution structure to insure that Vdd got to where it was needed (since it was already hooked up to Vdd). This effected the layout of power pads, lots of Vdd pads on the sides, lots of ground pads on the top and bottom.

The problem:

In the current designs the seal ring and crack protection ring (both 30 um wide) are separated by only 0.5 um. This is causing problems in the fab so the separation will be widened out to around 4 um on Euterpe and Mnemo but it can't be widened much more than this due to other fab limitations. Our air bridging scheme means that this gap will be unpassivated and it is, by definision, outside the seal ring. The presents of water vapor and chloride ions in the enviroments and the voltage difference between the seal and crack protection rings means that dendrites of gold will grow between the rings and short hte Vdd and ground power rails. This growth can occur very rapidly.

The solution:

The process and chip layout do not allow sufficient space between the seal and crack protection rings (~40 um are needed). The gap must, by definision, be outside the hermetic portion of the device where water and chloride ions can not be avoided. The crack protection ring will be shorted to the bulk of the die (ground) by the sawing operation. As a result the only solution I can think of is that the seal ring must be connected to ground, not Vdd. This will mean that the power distribution (at the pad ring level, not further in) will have to be redesigned with the assumption that the seal ring is ground and not Vdd. I would expect changes will include swapping Vdd and ground pins and changing the relative widths of poser and ground busses just inside the die pad ring.

I believe that this problem needs to be addressed now, before Euterpe or Mnemo tapes out and that it will effect several areas of the design to a greater or lesser extent.

Paul.

From: tbr
Sent: Monday, April 24, 1995 10:09 PM
To: steve
Subject: forwarded message from Ken Hsieh

Steve,

Background from SGI. As you see they recon "on site" response is 8hrs, an from what Scott said to me on the phone, that's buisness hours, which of course translates to 24hrs wall clock time. Ken also made the mistake of calling the engineer directly, rather than the hot-line, so the call did not get logged right away. I'll make sure in future he goes to the hot line first. Ken's call went in around 9:30 and the engineer was on site by about 3. So under the circumstances well within the 8 hours.

Score so far is that when the engineer took it down to install software patches the system disk failed (apparently a known problem on the IBM drives it uses). He went back to base to get a replacement drive, then when he tried to rebuild an OS on it he found he had a bad tape. He left around 6:30 and is due back in the morning. I have very low confidence the software patches are pertinent to the original problem.

Tim

----- Start of forwarded message -----
Return-Path: <ken>
Received: from rimulac.microunity.com by gaea.microunity.com (4.1/musel.3)
id AA04679; Mon, 24 Apr 95 14:47:11 PDT
Received: by rimulac.microunity.com (8.6.10/muse-sw.3)
id OAA07326; Mon, 24 Apr 1995 14:47:09 -0700
Message-ID: <199504242147.OAA07326@rimulac.microunity.com>
From: ken (Ken Hsieh)
To: scottm@maui.corp.sgi.com
Cc: tbr
Subject: Re: dump header file
Date: Mon, 24 Apr 1995 14:47:09 -0700

Scott,

Thanks a lot !
I will forward this to my manager, Tim Robinson, so we all know this.

Ken

> From scottm@maui.corp.sgi.com Mon Apr 24 14:42:12 1995
> From: "Scott Machtmes" <scottm@maui.corp.sgi.com>
> Date: Mon, 24 Apr 1995 14:41:50 -0700
> X-Mailer: Z-Mail (3.2.0 26oct94 MediaMail)
> To: ken@microunity.com (Ken Hsieh)
> Subject: Re: dump header file
> Mime-Version: 1.0
> Content-Type: > text/plain> ; > charset=us-ascii>
> Content-Length: 1086
>
> Ken,
> I will be out there about 3pm - 3:15pm.
>
> Please note a couple to things for clarification:
>
> 1. Your SGI support contract does not specify 4 hour on site response.
> Your
> standard full support is a 4 hour response with 8 hour on site response.

>
> 2. Please keep in mind that these responses are based upon the opening
> of a call to the SGI Hotline. If you call and leave me a message, I
> cannot insure that I will even hear it the same day.
>
> I am willing to do everything I can to accomodate you but you have to
> let me know what the current situation is and your current level of
> urgency for me to set both our expectations properly.
>
> Thanks,
>
> --
>
> -----
> Scott Machtmes, System Support Engineer, Silicon Graphics Computer Systems
> Internet : scottm@sgi.com
> UUCP : {sun,decwrl,pyramid,ucbvax}!sgi!scottm
> Telephone : (415) 390-3913
> US Mail : 2171 Landings Drive, Mountain View, CA 94043
> FAX : (415) 960-3391
> SGI office mail : Mail Stop DWR-275
>
> ----- End of forwarded message -----

From: graham (Graham Y. Mostyn)
Sent: Monday, April 24, 1995 8:46 PM
To: dbulfer; tbr; tbe@microunity.com
Cc: Philip; yves; graham
Subject: Re: Pandora module mechanical design meeting

Tom, I'm attending a course all this week, so unfortunately I won't be able to join you at the 3pm Wed review.

However, Jean-Yves and I call in at the office after class, so we could catch up with you then. Jean-Yves, perhaps you could discuss the approaches with Tom prior to the Wed meeting?

Graham.

> From tbe@microunity.com Mon Apr 24 16:06:58 1995
> Date: Mon, 24 Apr 95 16:06:55 PDT
> X-Sender: tbe@gaea.microunity.com
> Mime-Version: 1.0
> Content-Type: > text/plain> ; > charset="us-ascii">
> To: dbulfer, tbr, graham
> From: tbe@microunity.com (Tom Eich)
> Subject: Pandora module mechanical design meeting
> Cc: Philip
> Content-Length: 1040

> Hi,

> I would like to meet with you to review mechanical details of the common
> feature set to be used in the Hermes modules. These details have an
> impact on the pcb layouts of the four module types (Euterpe, Mnemo,
> Calliop, PCI/Hermes bridge), and also on the backplane. The Euterpe
> and Mnemo layouts have been completed to the level where the pcb
> routing constraints are well understood, and so we can now evaluate
> the mechanical design options that will accomodate these layout constraints.

> Can we meet on Wednesday, at 3:00pm in the boxers conference room?
> Because I would like to do some real-time concept selection (from
> design approaches I'll have available), I would like to limit the
> attendance, but if you want anyone else to be there, please let me know.

> -Tom

>
> Tom Eich
> MicroUnity Systems Engineering, Inc.
> 255 Caspian Dr. Sunnyvale, CA 94089
> (408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: Tom Eich [tbe@microunity.com]
Sent: Monday, April 24, 1995 6:07 PM
To: dbuffer; tbr; graham
Cc: Philip
Subject: Pandora module mechanical design meeting

Hi,

I would like to meet with you to review mechanical details of the common feature set to be used in the Hermes modules. These details have an impact on the pcb layouts of the four module types (Euterpe, Mnemo, Calliop, PCI/Hermes bridge), and also on the backplane. The Euterpe and Mnemo layouts have been completed to the level where the pcb routing constraints are well understood, and so we can now evaluate the mechanical design options that will accomodate these layout constraints.

Can we meet on Wednesday, at 3:00pm in the boxers conference room? Because I would like to do some real-time concept selection (from design approaches I'll have available), I would like to limit the attendance, but if you want anyone else to be there, please let me know.

-Tom

Tom Eich
MicroUnity Systems Engineering, Inc.
255 Caspian Dr. Sunnyvale, CA 94089
(408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: graham (Graham Y. Mostyn)
Sent: Friday, April 21, 1995 3:18 PM
To: geert; mudge; tbr; paulp
Cc: al; ahn; mouss; rich; graham
Subject: Re: processing of Castor/Pollux and Calliope1

Following a discussion with Paul, I would propose that after the new metal wafflization algorithms have been developed, and Euterpe/Mnemo has taken priority in tape-out, we turn our attention to implementing necessary metal changes on the Castor/Pollux mask set. Pollux is essential, among other things, for analysing analog structures on Euterpe (PLL, knob city etc), Calliope and Mnemo.

(We have started examining the Pollux PLL with a view to reducing differences between that older design and the Euterpe version).

Graham.

> From paulp Fri Apr 21 12:37:00 1995
> From: paulp (Paul Poenisch)
> Subject: processing of Castor/Pollux and Calliope1
> To: graham (Graham Y. Mostyn), geert (Geert Rosseel), mudge (john mudge),
> tbr (Tim B. Robinson)
> Date: Fri, 21 Apr 95 12:36:55 PDT
> Cc: al (Albert Matthews), ahn
> X-Mailer: ELM [version 2.3 PL11]
> Content-Length: 2146
>
> After the daily fab meeting yesterday afternoon it became apparent
> that we have some significant problems in processing Castor/Pollux and
> Calliope1 in the fab. The problems are different for the two reticle
> sets and are outlined below.
>
> Castor/Pollux:
>
> This reticle set was the first one that we designed. The design
> rules, fracture flow and wafflization methodologies have changed
> significantly since this reticle set was made. It now appears that
> the current reticle set can not be processed further than metal 1.
> This is due to serious metal peeling and dielectric delamination
> problems that occur because of the wafflization and perforation patterns that were used
> in the design of this device (mainly).
>
> As a result we have halted processing of this device past metal 1. We
> will still be running this device for transistor and contact
> characterization but Calliope-0, Pollux and the Castor tests dealing
> with layers above metal 1 are unusable at this time.
>
> Please note that this is not a case of the devices not working, the
> wafers can not be processed without risking contamination of the fab
> equipment, the fab and other lots with gold flakes.
>
> If the structures on this reticle set are needed for the business plan
> then it will need to be redone.
>
> Calliope-1
>
> The data used to generate this reticle set is not quite as upto date
> as Orchis (which is also not up to current standards) but it should be
> usable. Unfortunately when we processed the first lots of this
> device through the middle metal layers (metal 2 through 3) we
> discovered that the reticle vendor that made this set failed to hold
> the dimensional tolerance needed for this process. As a result we are having problems

processing this device.

>
> We believe that we can use the current reticle set to produce some
> Calliope-1's but the photomasking engineers will have to hand carry
> the lots through the metal layers. This will result in some slowing
> of the lots when compared to Orchis (which does not have this problem)
> but should still allow us to get enough devices out to allow debug of the device.
>
>
> I will keep you posted in changes to this situation.
>
> Paul.
>
>

From: wingard (Drew Wingard)
Sent: Friday, April 21, 1995 3:15 PM
To: solo
Cc: brian; fwo; geert; lisar; ong; tau; wampler
Subject: Re: missing layouts

Solo wrote:

```
> as Drew Wingard was saying .....
> ..
> ..Solo wrote:
> ..> there are a slew of missing layouts. any one want to figure out why?
> ...
> ..._MISSING_LAYOUT_FILE_ xs2andn10s.ly
> .. <snip>
> ... solo a.k.a. John Campbell x516
> ..
> ...I see them in /u/chip/atlas/compass/leaf. Are we not using the
> proper ..vlsi.boo file somewhere?
> ..
> ..(These are the actual standard cells layouts, as produced by the
> atlas leafmold) ..
> ..Drew
> ..
> ..
> ./compass/vlsi.boo-all
> ./compass/vlsi.boo-dcell
> ./compass/vlsi.boo-tapeout
> ....
> regards,
> solo a.k.a. John Campbell x516
```

and then Solo wrote:

```
> ok so why does Depend-pdl think they are not there
```

Looks like the atlasleaf.db file in atlas/compass/leaf needs updating (i.e. looks like vlsifixlib need to be run).

Possibly a bug in atlas/leafgen/leafmold's Makefile. Tom???

Drew

From: solo (John Campbell)
Sent: Friday, April 21, 1995 3:13 PM
To: solo
Cc: wingard@echidna.microunity.com; brianl (Brian Lee); fwo (Fred Obermeier); geert (Geert Rosseel); lisar (Lisa Robinson); ong (Warren R. Ong); tau; wampler (Kurt Wampler)
Subject: Re: missing layouts

as solo was saying
..
..as Drew Wingard was saying
....
....Solo wrote:
....> there are a slew of missing layouts. any one want to figure out why?
....>
....> _MISSING_LAYOUT_FILE_ xs2andn10s.ly
.... <snip>
....> solo a.k.a. John Campbell x516
....
....I see them in /u/chip/atlas/compass/leaf. Are we not using the propervlsl.boo
file somewhere?
....
....(These are the actual standard cells layouts, as produced by the atlas leafmold)
....Drew
....
..
.../compass/vlsl.boo-all
.../compass/vlsl.boo-dcell
.../compass/vlsl.boo-tapeout
..
ok so why does Depend-pdl think they are not there.

....
regards,
solo a.k.a. John Campbell x516

From: solo (John Campbell)
Sent: Friday, April 21, 1995 3:10 PM
To: Drew Wingard
Cc: briant (Brian Lee); fwo (Fred Obermeier); geert (Geert Rosseel); lisar (Lisa Robinson); ong (Warren R. Ong); tau; wampler (Kurt Wampler)
Subject: Re: missing layouts

as Drew Wingard was saying

..
..Solo wrote:
..> there are a slew of missing layouts. any one want to figure out why?
..>
..> _MISSING_LAYOUT_FILE_ xs2andn10s.ly
..> <snip>
..> solo a.k.a. John Campbell x516
..
..I see them in /u/chip/atlas/compass/leaf. Are we not using the proper ..vlsi.boo file
somewhere?
..
..(These are the actual standard cells layouts, as produced by the atlas leafmold) ..
..Drew
..
..
./compass/vlsi.boo-all
./compass/vlsi.boo-dcell
./compass/vlsi.boo-tapeout
..
regards,
solo a.k.a. John Campbell x516

From: geert (Geert Rosseel)
Sent: Friday, April 21, 1995 1:21 PM
To: bpw; ong; stick; vo; wingard
Cc: bill; michael; tbr
Subject: Re: TSMC foundry models

> Geert - what are our marching orders here?

We should have a design that we can tape-out to both foundries. Dave has a set of the design rules and he'll check if they are compatible. They are supposed to be fully compatible, so unless we see any surprises, we should be O.K. on that.

In terms of SPICE simulations, it is not clear to me how different the models are. I suspect they are quite similar. If so, I suggest we stay with CSM as our standard simulation environment and verify critical cells with TSMC models to make sure they still work. Bruce, B,P, : can you play around with the models and find out how different they are ?

Geert

From: geert (Geert Rosseel)
Sent: Thursday, April 20, 1995 11:43 AM
To: bill; billz; bpw; dickson; stick; tbr; vo; woody
Subject: Euterpe custom block interface reviews

Hi,

We need to do one last review on the timing of all interfaces to the custom blocks on Euterpe. Let's meet at 3:00 p.m. to discuss what needs to be done and how.

Hardware Conference Room
Thursday 3:00 p.m.

Geert

From: lisar (Lisa Robinson)
Sent: Monday, April 17, 1995 1:51 PM
To: solo
Cc: tbr; ~/Mail/euterpe; geert
Subject: Re: Release of BOMs by tbr (proteus)

> From solo Mon Apr 17 11:37:29 1995
> Return-Path: <solo>
> Received: from echidna.microunity.com by gaea.microunity.com (4.1/musel.3)
> id AA17125; Mon, 17 Apr 95 11:37:29 PDT
> Received: by echidna.microunity.com (8.6.10/muse-sw.3)
> id LAA08061; Mon, 17 Apr 1995 11:37:28 -0700
> From: solo (John Campbell)
> Message-Id: <199504171837.LAA08061@echidna.microunity.com>
> Subject: Re: Release of BOMs by tbr (proteus)
> To: tbr@microunity.com (Tim B. Robinson)
> Date: Mon, 17 Apr 95 11:37:27 PDT
> Cc: geert (Geert Rosseel), lisar (Lisa Robinson)
> In-Reply-To: <199504171827.LAA00399@aphrodite.microunity.com>; from
> "Tim B. Robinson" at Apr 17, 95 11:27 am
> X-Mailer: ELM [version 2.3 PL11]
> Content-Length: 2484
> X-Lines: 65
> Status: RO

> as Tim B. Robinson was saying

> ..

> ..

> ..John Campbell wrote (on Mon Apr 17):

> ..

> .. as Tim B. Robinson was saying

> ..

>Checkin Message: -----

>added dlatcqtq.V

> ..

>BOM Update in proteus BOM 5.1157

>BOM Update in proteus/verilog BOM 5.149

>BOM Release in proteus/verilog/libsrc BOM 147.0

> ..

> ..

> ..

> .. to pick up the latest, (since i am stalled) do i just do a getbom from
> .. the top or do i need to update BOM and then getbom.

> ..

>regards,

> .. solo a.k.a. John Campbell x516

> ..

> ..If you want to get this into the snapshot, you should getbom from
> the ..top having checked to see what else will be picked up. However,
> we ..need to be sure geert is willing to have the make re-run at this
> ..point. I assume you are stalled for something else, not this?

> ..

> ..Tim

> ..

At some point (prior to tape-out) I will need this in the snapshot.

Lisa R.


```

> ..
>
> yes.  the license for iss lpe and such expired and stalled the build.
>
> the following get picked up
> not very illustrative since it doesn't tell you what directory.
> Dir      proteus/verilog/libsrc          BOM 147.0
> 1.191    Makefile                        (1.190)
>
> 1.2      dlatchq.V                       (1.1)
> 146.1    dlatchtq.V                      (No)
> 1.4      serbiflop.V                     (1.3)
> 1.4      sertriflop.V                    (1.3)
> 1.4      xclatbc.V                       (1.3)
> 1.2      xcloadlatbc.V                   (1.1)
> 1.4      xcnrlatbc.V                     (1.3)
> 2.1      .checkoutrc                     (No)
>
> Dir      proteus/verilog/src/dram        BOM 3.0
> 1.1      Makefile                        (No)
> 1.1      dram.V                          (No)
>
> Dir      proteus/verilog/zblibsrc       BOM 69.0
> 1.27     Makefile                        (1.26)
>
> 8.28     cerbsnoop.c                     (8.27)
> 64.3     cerbsnoop_events.h              (64.2)
> 9.24     he.c                             (9.23)
> 1.1      ttlc2emn.v
>
> ....
> regards,
> solo a.k.a. John Campbell   x516
>
>
>

```

From: wingard (Drew Wingard)
Sent: Monday, April 17, 1995 11:51 AM
To: hardheads; softheads
Subject: This week on SITN

This first one is probably most interesting to hardware types, but the second one is definitely software/test material.

Regards,
Drew

My SITN schedule sez that EE310 is tape-delayed on E2 from 7:15-8:30 on Tues.
SITN info for EE380 is in the posting.

***** begin included messages *****

EE310 Seminar: April 18, McCullough 134, 4:15 pm

IBM PowerPC Technology Overview

David C. Thomas
E.J. Nowak

IBM Microelectronics Division
Essex Jct., VT

ABSTRACT:

In 1991 the PowerPC alliance was formed by Apple, IBM and Motorola to deliver a family of high-performance, scalable RISC microprocessors. Fueling the continued drive to higher performance and density is IBM's CMOS technology, developed by the IBM Microelectronics Division. Beginning with the first PowerPC 601 in 0.8/0.6um 3.6V CMOS, a steady progression of technologies are discussed. Innovations include shallow trench isolation, low-voltage quarter-micron MOSFETs, chemical-mechanical polishing, damascene local interconnects and multi-level metalization. These and other advances currently being manufactured in the Essex Junction fabricators are reviewed. Finally, we explore future challenges to further CMOS technology scaling, with particular focus on the increasing importance of delays associated with interconnect RC.

BIOGRAPHY: David C. Thomas

David C. Thomas received his B. Sc. degree in electrical engineering from Wilkes College in 1984. After an internship at the IBM Thomas J. Watson Research Center in Yorktown Heights, NY, where he helped to demonstrate the first experimental evidence of ballistic transport in GaAs/AlGaAs materials, he received his masters of engineering degree from Cornell University in 1986. His work with Professor Simon Wong in the use of selective Tungsten for multilevel interconnects earned him a Ph.D. in electrical engineering from Cornell University in 1990. He joined IBM Microelectronics in Burlington, VT, as a staff level engineer in 1990 and was promoted to advisory engineer in 1994. His current position involves the development, implementation and qualification of advanced CMOS logic technologies into a 200mm manufacturing line. Dr. Thomas has authored eight patents and ten technical papers, and was recently appointed to the interconnect subcommittee of the IEDM.

EE380 Computer Systems Colloquium

Spring Quarter 1994-1995

Lecture #3

Date: Wednesday, Apr 19, 1994
Time: 4:15-5:30 pm
Location: Terman Auditorium
SITN: Thursday, Channel E3, 8:00-9:15
Speaker: Barton P. Miller
Computer Sciences Department
University of Wisconsin
Title: Making Programs Explode: Using Simple Random
Testing on Real Programs

Abstract

In 1990, we published the results of a study of the reliability of standard UNIX utility programs. This study showed that by using simple (almost simplistic) random testing techniques, we could crash or hang 25-33% of these utility programs. Five years later, we have repeated and significantly extended this study using the same basic techniques: subjecting programs to random input streams. A distressingly large number of UNIX utilities still crash with our tests.

We tested a wide variety of utility programs on nine UNIX platforms. The programs were sent random input streams. We used a conservative and crude measure of reliability: a program is considered unreliable if it crashes with a core dump or hangs (infinite loop). We used the random testing to also test X-Window applications and servers, network servers, and system library interfaces.

The major results of this study are: (1) In the last five years, all previously-tested versions of UNIX made noticeable improvements in the reliability of their utilities. But ... the failure rate of these systems is still distressingly high (from 18-23% in the 1995 study). (3) Even worse is that many of the same bugs that we reported in 1990 are still present in the code releases of 1995. (4) The failure rate of utilities on the commercial versions of UNIX that we tested (from Sun, IBM, SGI, DEC, and NEXT) ranged from 15 to 43%. (5) The failure rate of the utilities on the freely-distributed Linux version of UNIX was second-lowest, at 11%. (6) The failure rate of the public GNU utilities was the lowest in our study, at only 7%. (7) We could not crash network services on any of the versions of UNIX that we tested. (8) Almost a quarter of the X-Window applications that we tested crash on purely random input data streams (random binary data). More significant is that more than 40% of the applications crash given random, but legal X-event streams. (9) We could not crash X server on the versions of UNIX that we tested (i.e., sending random data streams to the server).

Biography

Bart Miller joined the Computer Sciences Department at the University of Wisconsin-Madison in 1984, where he is currently a Professor. He received his M.S. and Ph.D. in Computer Science from Berkeley in 1980 and 1984. His current research interests include parallel programming tools, network name services, and mobile computing.

Contact Information:

Barton P. Miller
Computer Sciences Department

University of Wisconsin
1210 W. Dayton
Madison, WI 53706-1685
bart@cs.wisc.edu
bart@wis.stanford.edu

```
*****
* EE380 is the Computer Systems Laboratory Colloquium.  The Colloquium *
* meets most Wednesdays throughout the academic year.                *
* *                                                                      *
* LECTURES ARE OPEN TO EVERYONE -- FACULTY, STUDENT (ENROLLED OR NOT) *
* INDUSTRIAL VISITORS, OR OTHER INTERESTED PARTIES.                    *
* *                                                                      *
* We frequently have a "dutch treat" dinner for the speaker following *
* the lecture.  If you would like to join us, please contact one of    *
* the course organizers.                                                *
* *                                                                      *
* For information on the class send e-mail with a subject line          *
* mentioning "info" in the subject line to ee380@shasta.stanford.edu.  *
*****
```

From: trancy [trancy@charybdis]
Sent: Monday, April 17, 1995 12:06 PM
To: Geert Rosseel
Subject: RE: Summary of Cronus Meeting

Morning Geert,

Since testing is another critical area for the success of Cronus, I think Johnny should be in the mailing list and attend the meeting.

Regards.
Trancy.

From: Geert Rosseel on Sun, Apr 16, 1995 10:10 AM
Subject: Summary of Cronus Meeting
To: bill; bpw; brianl; dane; fwo; geert; hopper; lisar; ong; paulp; solo; stick; tbr; tom; trancy; vanthof; vo; wampler; wingard

Hi,

Here is a summary of what we've discussed in Friday's Cronus meeting. I took the liberty of adding the names of the persons responsible on the different sections.

Not entered in this are verification and required logic design changes. Logic changes would have to happen in May and verification before July 15.

TAPE-OUT DATE : July 15

Baseplate :

* floorplan	Warren
* clock spars	Bill (Design), Warren (Baseplate Makefile), Kurt (clock program)
* power distribution	Bill (Design)
* padding	Warren (Baseplate Makefile), Warren

Deadline :

May 1	: Have a GARDS Compilable baseplate
June 1	: Finished Baseplate (exact die size and pad-ring) Working automated clock spar generation.

Custom Blocks :

* Cache	Bruce
* Tag	Bruce
* TLB	BP
* Register File	BP
* NB	BP
* TTL I/O blocks	Bruce, BP
* Hermes Channel	Dane, Bill

Deadline :

June 1	: All custom blocks layed out
--------	-------------------------------

Standard Cells :

* Schematics	Warren
--------------	--------

* Verilog	Brianl
* Timing	Fred
* Layout	
XL	Warren
XS	Tom (L.)
* PDL	Tom (L.)

Deadline :

May 1	: Have a GARDS Compilable complete Atlas library
	Solo to build the Atlas database
June 1	: Atlas database finished

Mapping / Place and Route

* Methodology	Brianl, Drew
* Implementation	ALL

Deadline :

May 1	: Have an initial working Makefile.rules checked in
June 1	: Makefile.rules fully functional
July 1	: Fished place and route of a functional Cronus

Packaging :

* MCM	
* Tab	Trancy
* "Simple packages"	

Deadline :

April 21	: Decide on strategy
Other dates depend on the packaging scheme that we decide on.	

From: geert (Geert Rosseel)
Sent: Sunday, April 16, 1995 12:10 PM
To: bill; bpw; briani; dane; fwo; geert; hopper; lisar; ong; paulp; solo; stick; tbr; tom; trancy; vanthof; vo; wampler; wingard
Subject: Summary of Cronus Meeting

Hi,

Here is a summary of what we've discussed in Friday's Cronus meeting. I took the liberty of adding the names of the persons responsible on the different sections.

Not entered in this are verification and required logic design changes. Logic changes would have to happen in May and verification before July 15.

TAPE-OUT DATE : July 15

Baseplate :

* floorplan	Warren
* clock spars	Bill (Design), Warren (Baseplate Makefile), Kurt (clock program)
* power distribution	Bill (Design)
* padding	Warren (Baseplate Makefile), Warren

Deadline :

May 1	: Have a GARDS Compilable baseplate
June 1	: Finished Baseplate (exact die size and pad-ring) Working automated clock spar generation.

Custom Blocks :

* Cache	Bruce
* Tag	Bruce
* TLB	BP
* Register File	BP
* NB	BP
* TTL I/O blocks	Bruce, BP
* Hermes Channel	Dane, Bill

Deadline :

June 1	: All custom blocks layed out
--------	-------------------------------

Standard Cells :

* Schematics	Warren
* Verilog	Brianl
* Timing	Fred
* Layout	
XL	Warren
XS	Tom (L.)
* PDL	Tom (L.)

Deadline :

May 1	: Have a GARDS Compilable complete Atlas library Solo to build the Atlas database
June 1	: Atlas database finished

Mapping / Place and Route

- * Methodology Brianl, Drew
- * Implementation ALL

Deadline :

May 1 : Have an initial working Makefile.rules checked in
June 1 : Makefile.rules fully functional
July 1 : Fished place and route of a functional Cronus

Packaging :

- * MCM
- * Tab Trancy
- * "Simple packages"

Deadline :

April 21 : Decide on strategy
Other dates depend on the packaging scheme that we decide on.

From: Tom Eich [tbe@microunity.com]
Sent: Saturday, April 15, 1995 7:20 PM
To: tbr
Subject: problem with the DTM's controller

Hi,

Kien and I had a problem with the DTM (sintering machine) when we were preparing to run the Pandora Euterpe module for the board meeting. Ken was here by the time we realized that we needed sysadmin help, as the PC that controls the DTM was asking for a login and password that we don't know (JR's didn't work). Ken has been working on it for over an hour, but doesn't seem to be making much progress (ericm had set up this oddball system [pc running unix]).

I don't feel it's helpful for me to push on ken, as he's doing his best with limited knowledge of this machine, but we can't get a sinter run going without getting the machine to give us an X window. Any help or prioritization you can provide is appreciated. Its looking iffy for a model if we don't get this resolved by Monday am, and I have to leave 10 minutes ago.

Thanks,

-Tom

Tom Eich
MicroUnity Systems Engineering, Inc.
255 Caspian Dr. Sunnyvale, CA 94089
(408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: pmayer (Patricia Mayer)
Sent: Thursday, April 13, 1995 6:39 PM
To: tbr
Cc: pmayer
Subject: Re: Pandora Euterpe PCB review

> From: tbr Thu Apr 13 16:28:12 1995
> Date: Thu, 13 Apr 1995 16:28:07 -0700
> From: tbr (Tim B. Robinson)
> To: pmayer (Patricia Mayer)
> Cc: dbulfer, howard, lisar, pandora, philip, pmayer, woody
> Subject: Pandora Euterpe PCB review
> Content-Length: 440
>
>
> Patricia Mayer wrote (on Thu Apr 13):
>
> We will have a final Pandoa-Euterpe PCB layout review tomorrow,
> Friday the 14 at 10:00 in the Engineering conference room. Please let me
> know of any scheduling conflicts.
>
> I have a meeting with a vendor at 9.30 I may done by 10, but it may
> be 10.30, so could we delay till then please?
>

Tim,

Tom won't be in Friday so this was postponed until Monday the 17th at 10:00. Please let me know if you didn't get mail on this or if this is a new conflict for you.

-Pattie

From: tbr (Tim B. Robinson)
Sent: Thursday, April 13, 1995 6:28 PM
To: pmayer (Patricia Mayer)
Cc: dbulfer; howard; lisar; pandora; philip; pmayer; woody
Subject: Pandora Euterpe PCB review

Patricia Mayer wrote (on Thu Apr 13):

We will have a final Pandoa-Euterpe PCB layout review tomorrow,
Friday the 14 at 10:00 in the Engineering conference room. Please let me
know of any scheduling conflicts.

I have a meeting with a vendor at 9.30 I may done by 10, but it may be 10.30, so could we
delay till then please?

Remember we are only going out to film. We might plan to have a film
review once these come in.

From: pmayer (Patricia Mayer)
Sent: Thursday, April 13, 1995 1:31 AM
To: tbr
Subject: Re: Board plots

> From: tbr Wed Apr 12 22:29:04 1995
> Date: Wed, 12 Apr 1995 22:29:02 -0700
> From: tbr (Tim B. Robinson)
> To: pmayer
> Subject: Board plots
> Content-Length: 154
>
>

> There will be a boards meeting tuesday. Mouss has asked for plots of
> the Pandora Euterpe and Mnemo modules. Can you arrange to make them
> please?
>
>

> Tim
>
>

Sure Tim.

I assume I will still be alive after they've installed my ISDN line. Ha ha.

-Pattie

From: tbr
Sent: Thursday, April 13, 1995 12:29 AM
To: pmayer
Subject: Board plots

There will be a boards meeting tuesday. Mouss has asked for plots of the Pandora Euterpe and Mnemo modules. Can you arrange to make them please?

Tim

From: tbr
Sent: Wednesday, April 12, 1995 2:14 AM
To: Curtis Abbott
Cc: agc; tbr
Subject: followup

Curtis Abbott wrote (on Thu Apr 6):

Here's a summary of our discussion with some questions. Perhaps you can respond to questions and mistakes. My thought is to send a version of this to you, Scott and Alan as talking points for a meeting or meetings aimed at proposing directions for a cost-effective settop box (ie. mpeg, qam, ntsc).

- Curtis

<intro would go here. need for focus on system level, i.e. both Eu and Ca, and on power & cost. also, desire to leverage design work we've already done.>

Some ideas about Euterpe:

Step 1. go to knob setting 1

- power = 17W
- speed = 250 MHz (?)

Since we are currently at knob 6, going to 1 with direct scaling of both power and speed would give 14W and 18MHz before optimization.

Step 2. timing optimize for knob setting 1

- no power change
- speed = 330 MHz (?)

It looks like 300MHz will be hard to reach at knob 1. Knob 2 of course double power, however, for a given cycle time, the power optimizer will power down so we gain some back. My guess though is it will only take it down 20% before essentially everything is a min power cell where we again get a sub-optimal design. We have 2 new ideas here. First add some new even lower power cells, second use the "process code" to get the effect of a knob setting 1.5. Some combination of these may let us get 300MHz at order 20W.

Step 3. change leaf cell family for knob setting 1

- atom count goes down; 3% of total die area (???)
- fewer parasitics so heavily powered nets get faster (what's the bottom line?)

Step 4. change atom for knob setting 1

- remove bias xtors
- other changes?
- atom gets smaller; 3% of total die area (???)

I don't think this one affects die area or utilization unless we actually change atom dimensions. That's more than we'd want to do, since all the custom structures have interfaces designed to mesh into the existing standard sofa grid.

Step 5. SRAM changes for lower speed

- cmos read ports
- cache/buffer loads issue 1/slot; stores 0.5/slot
- reduce power (12/5 W -> ???)
- better software IPC

Step 6. remove 1 Hermes channel

- regain 4% of total die area?

Step 7. rework NB/DRAM interfaces for knob setting 1
- should reduce area, give better (relative) latency

I don't remember what we were proposing here. I see the problem as being that the existing design has problems when the SDRAM clock ratio is set below 10. Fixing this is not trivial and unlikely to save area.

Step 8. add specialized HW and/or new instructions
- e.g. huffman decoder, averaging, rounding, sine wave generation
- goal would be to reduce required instructions for settop benchmark by factor of 2, while still fitting in 1 cm²

Further thoughts after a further discussion with alan.

If the cable modem can be built without requiring *any* DRAM, then the caches seem to make little sense (nothing to back them from), so I assume you'd have to set the cache size to 0. If so we are paying a big price for tags, and gtlb for something which is almost totally running out of on chip memory. We might free up a lot of area there if we were willing to simplify memory management and eliminate caches entirely in favor of just the buffers. Of course this would have major implications for design time and verification, not to mention that it would be architecturally incompatible with euterpe.

Tim

From: wayne@microunity.com
Sent: Monday, April 10, 1995 2:15 PM
To: lisar@rhea
Cc: dbulfer@rhea; doi@rhea; graham@rhea; guarino@rhea; philip@rhea; tbr@rhea
Subject: ECR/2192: Main board rev change from 620-00001-0000 rev2 to ??

>Number: 2192
>Category: ECR
>Synopsis: Main board rev change from 620-00001-0000 rev2 to ??
>Confidential: yes
>Severity: critical
>Priority: high
>Responsible: lisar (Lisa Robinson)
>State: open
>Class: change-request
>Submitter-Id: MUSE
>Arrival-Date: Mon Apr 10 12:15:00 1995
>Originator: Wayne Freitas
>Organization: MicroUnity Systems Engineering, Inc.
>Release: 620-00001-0000 rev2
>Environment:

>Description:
This is the ECR request to change the main board PCB (620-00001-0000 rev2) to a new revision. The change requirement is to allow continued test development efforts to meet Hestia technical requirements.

No additional features are being added to this revision.
No software features apply at this time to this PCB.
PR's used:

#	Description	
1771	Decoupling caps connect to wrong plane	
1772	Surface mount caps change to through hole	
1774	UL violation	DRC issue
1777	IR receivers ground pin connected to wrong gnd	
1792	Fan return needs isolated gnd	
1806	Smart card power vias to small	
1809	Add fuse to fan	
1825	Dead traces on layers 1 & 3	DRC issue
1826	Soldermask issue, causes copper exposure	DRC issue
1827	Soldermask strips below .005"	DRC issue
1828	PTH on center of TAB error	DRC issue
1829	PTH diameter clearance issue	DRC issue
1830	Aspect ratio	DRC issue
1831	Breakaway holes clearance to copper layer	DRC issue
1860	Floating ground planes	DRC issue
1861	Pinout error on DC-DC Module	
1862	Fabrication error	DRC issue
1869	Duplicate reference designator	Tool
1876		
1877	Duplicate reference designators	Tool
1878	Mounting pad has incorrect hole diameter	
1879	370-00010-0000 (RCA) has incorrect foot print pattern	
1880	370-00021-0000 (mono jack) solder holes to large	
1881	430-00001-0000 (RF relay) has incorrect landing pads	
1882	Incorrect reference designator number	Tool
1901	Floating leads on multiple components (gnd missing)	
1902	Eleetrolytic caps have incorrect pinouts	

		Tool
1905	Duplicate reference designators	
1909	Floating leads on multiple components (gnd missing)	
1910	VCO noise issue (trace routing)	
1911	VCOs need additional PS filtering	
1914	DC-DC -sense line shorted	
1925	270-00004-0000 landing pattern incorrect	
1929	Schematic <--> BOM discrepancies	
1930	Wrong reference designator	
1931	Pad spacing incorrect on SDRAMs	
1932	Inadequate clearance on A30C1 and bottom of box	
1933	Inadequate clearance between DC-DC fasteners & traces	
1934	Wrong reference designator	
1945	Floating leads on multiple components (gnd missing)	
1946	Floating lead on component (gnd missing)	
1950	Short between analog and digital gnds	
1959	Insufficient solder pad under VCOs	
1960	Manufacturability of attaching VCOs to main board	
1962	Inadequate clearance for TAB tooling under Euterge	
1979	Breakaway tabs	
1993	110-00007-0000 has incorrect pinout (QPSK amp)	
2001	Floating leads on Mini DIN (gnd missing)	
2004	Chassis plane extends inside phone barrier	
2007	Excessive noise on 3.3V power plane	
2009	Excessive noise on 5V power plane	
2010	Excessive noise on 12V power plane	
2012	Additional filtering needed on Audio/Video connectors	
2014	Trace rerouting	
2031	Hermes connector pinout changed	
2035	Baluns performance problem	
2036	AGC performance problem	
2059	Diplexer pinout problem	
2104	Hermes expansion connectors missing chassis gnd	
2105	Silkscreen outline on Flash ROM	
2179	Component change on Video input	

>How-To-Repeat:

na

>Fix:

>Audit-Trail:

>Unformatted:

From: tbr (Tim B. Robinson)
Sent: Sunday, April 09, 1995 10:54 AM
To: lisa
Cc: euterpe@news

Lisa Repka wrote (on Fri Apr 7):

In article <1995Apr6.192525.1147@microunity.com>, I wrote:
> I must be missing something -- can anybody clear up my confusion?

It seems there is some general interest over this issue (since several people have already asked me if I got a response to my previous posting), so I figured I'd share the answer:

> Other changes in implementation, while they have substantial performance
> effects, can be dealt with by using the Euterpe subset instruction set
> for prototyping purposes. This change to the VM system, if not implemented,
> would be difficult to work around. It is also a change that involves
> minimal hardware effects, and has been previously discussed with the
> hardware implementers.
>
> This hack, as you call it, can be broadly applied to any software application
> which requires > 47 bit addressing. It need not affect the operation of
> any current software, as the GA field can be configured to make the
> machine operate the same as previously.
>
> So, yes, Lisa, it was seriously intended for inclusion in the current
> Euterpe.
>
> Regards,
> Craig

I understand now. But to avoid my running into similar confusion in the future, I just have one more question -- When do we say our specification of the architecture for this -- our first -- revision of the cpu is finished? I know the standard cute reply is "when we tape it out". I even know that there is some truth to that, but I always thought it meant that the tape is the "ultimate documentation", which might (unintentionally) differ from other, separate, documentation. But in order to complete at least minimal verification before taping out, I thought that there had to be some point after which no changes in specification/implementation are allowed...

[Please excuse my ignorance; I clearly have jumped to some conclusions based on my own assumptions, rather than on any MU policies/schedules. The only product experience I have is in software, where we had freeze dates, after which fixes for severe bugs were the only changes that could be made, with NO exceptions, and any change was followed by a full, extremely time-consuming, test-cycle. I had no information that we had passed such a freeze-date for Euterpe, I had just, incorrectly, *assumed* that we had.]

I think the one thing which was wrong in this case is that the issue was not more widely discussed at the point where craig first raised it some time ago and discussed the implications with the designers. We have to seriously weigh the costs (time, delay, atoms etc) against the benefits for any change before deciding to make it. In this case, (unlike the addition of average instructions, which we said we could not include in this version), the costs are minimal, as are the risks, since as craig points out, by default we can make the behaviour identical to before. We did also discuss the implications with the verification group to make sure the issues there would not be significant. Lisar's input was that while additional test cases would be needed to cover the new functionality, there would be no effect on existing tests and that the additional cases

could be provided without needing this feature to be supported in terp first (we have a lot of other cases where tests do not run on terp already).

It is important that we do not let gratuitous changes interfere with our critical path, but equally we must not lose the agility we have as a small group to make improvements even at a late stage. After all, we are not a large company with armies of contributors none of whom sees more than a tiny piece of the picture and where process rigidity is a requirement to get a job done at all.

To put this particular case in perspective, we have a formal release process which would protect us in the case that for some reason the change caused an unexpected complication (ie easy to back out).

Releases get made approximately daily (after discussions within the implementation group) with what will be included. In recent weeks we have been making bug fixes, placement improvements, routing improvements and timing fixes, many of which are more significant than the changes implied by this architecturally visible change. So, while from the outside it may look like this is a scary change to be including at this late point, it will get rolled in with other changes with minimal impact.

To address this issue that this should have been more widely discussed earlier, I will repost craig's mail as an 'IMMINENT DECISION'.

Tim

From: tbr
Sent: Sunday, April 09, 1995 10:54 AM
To: lisa
Cc: euterpe@news

Lisa Repka wrote (on Fri Apr 7):

In article <1995Apr6.192525.1147@microunity.com>, I wrote:
> I must be missing something -- can anybody clear up my confusion?

It seems there is some general interest over this issue (since several people have already asked me if I got a response to my previous posting), so I figured I'd share the answer:

> Other changes in implementation, while they have substantial performance
> effects, can be dealt with by using the Euterpe subset instruction set
> for prototyping purposes. This change to the VM system, if not implemented,
> would be difficult to work around. It is also a change that involves
> minimal hardware effects, and has been previously discussed with the
> hardware implementers.
>
> This hack, as you call it, can be broadly applied to any software application
> which requires > 47 bit addressing. It need not affect the operation of
> any current software, as the GA field can be configured to make the
> machine operate the same as previously.
>
> So, yes, Lisa, it was seriously intended for inclusion in the current
> Euterpe.
>
> Regards,
> Craig

I understand now. But to avoid my running into similar confusion in the future, I just have one more question -- When do we say our specification of the architecture for this -- our first -- revision of the cpu is finished? I know the standard cute reply is "when we tape it out". I even know that there is some truth to that, but I always thought it meant that the tape is the "ultimate documentation", which might (unintentionally) differ from other, separate, documentation. But in order to complete at least minimal verification before taping out, I thought that there had to be some point after which no changes in specification/implementation are allowed...

[Please excuse my ignorance; I clearly have jumped to some conclusions based on my own assumptions, rather than on any MU policies/schedules. The only product experience I have is in software, where we had freeze dates, after which fixes for severe bugs were the only changes that could be made, with NO exceptions, and any change was followed by a full, extremely time-consuming, test-cycle. I had no information that we had passed such a freeze-date for Euterpe, I had just, incorrectly, *assumed* that we had.]

I think the one thing which was wrong in this case is that the issue was not more widely discussed at the point where craig first raised it some time ago and discussed the implications with the designers. We have to seriously weigh the costs (time, delay, atoms etc) against the benefits for any change before deciding to make it. In this case, (unlike the addition of average instructions, which we said we could not include in this version), the costs are minimal, as are the risks, since as craig points out, by default we can make the behaviour identical to before. We did also discussed the implications with the verification group to make sure the issues there would not be significant. Lisar's input was that while additional test cases would be needed to cover the new functionality, there would be no effect on existing tests and that the additional cases

could be provided without needing this feature to be supported in terp first (we have a lot of other cases where tests do not run on terp already).

It is important that we do not let gratuitous changes interfere with our critical path, but equally we must not lose the agility we have as a small group to make improvements even at a late stage. After all, we are not a large company with armies of contributors none of whom sees more than a tiny piece of the picture and where process rigidity is a requirement to get a job done at all.

To put this particular case in perspective, we have a formal release process which would protect us in the case that for some reason the change caused an unexpected complication (ie easy to back out).

Releases get made approximately daily (after discussions within the implementation group) with what will be included. In recent weeks we have been making bug fixes, placement improvements, routing improvements and timing fixes, many of which are more significant than the changes implied by this architecturally visible change. So, while from the outside it may look like this is a scary change to be including at this late point, it will get rolled in with other changes with minimal impact.

To address this issue that this should have been more widely discussed earlier, I will repost craig's mail as an 'IMMINENT DECISION'.

Tim

From: lisa (Lisa Repka)
Sent: Friday, April 07, 1995 6:55 PM
To: euterpe@news

In article <1995Apr6.192525.1147@microunity.com>, I wrote:
|> I must be missing something -- can anybody clear up my confusion?

It seems there is some general interest over this issue (since several people have already asked me if I got a response to my previous posting), so I figured I'd share the answer:

> Other changes in implementation, while they have substantial
> performance effects, can be dealt with by using the Euterpe subset
> instruction set for prototyping purposes. This change to the VM
> system, if not implemented, would be difficult to work around. It is
> also a change that involves minimal hardware effects, and has been
> previously discussed with the hardware implementers.
>
> This hack, as you call it, can be broadly applied to any software
> application which requires > 47 bit addressing. It need not affect the
> operation of any current software, as the GA field can be configured
> to make the machine operate the same as previously.
>
> So, yes, Lisa, it was seriously intended for inclusion in the current
> Euterpe.
>
> Regards,
> Craig

I understand now. But to avoid my running into similar confusion in the future, I just have one more question -- When do we say our specification of the architecture for this -- our first -- revision of the cpu is finished?
I know the standard cute reply is "when we tape it out". I even know that there is some truth to that, but I always thought it meant that the tape is the "ultimate documentation", which might (unintentionally) differ from other, separate, documentation. But in order to complete at least minimal verification before taping out, I thought that there had to be some point after which no changes in specification/implementation are allowed...

[Please excuse my ignorance; I clearly have jumped to some conclusions based on my own assumptions, rather than on any MU policies/schedules. The only product experience I have is in software, where we had freeze dates, after which fixes for severe bugs were the only changes that could be made, with NO exceptions, and any change was followed by a full, extremely time-consuming, test-cycle. I had no information that we had passed such a freeze-date for Euterpe, I had just, incorrectly, *assumed* that we had.]

From: tbr
Sent: Friday, April 07, 1995 2:02 PM
To: Curtis Abbott
Subject: something I forgot to mention

Curtis Abbott wrote (on Fri Apr 7):

Tim B. Robinson wrote (on Fri Apr 7):

Curtis Abbott wrote (on Thu Apr 6):

FYI. Relative to your call on Craig's latest proposal... There's some sentiment beyond Lisa over here that changing the Euterpe spec now would indicate a lack of discipline on our part. I see Craig's proposal changing verification code and perhaps a little bit of our code as well as logic design and layout. I think the impact on the software side is small; my concern is more with the engineering process issues.

Well, this is something he actually raised with a few people including me a good time back but did not publish till we had looked at what it would take. It's actually trivial to implement, so it was hard to argue against including it. I agree that we don't demonstrate as much discipline as we should, but realistically, I don't see this having any impact on tapeout, given we still have bug fixes which must be included.

When you say it's trivial to implement, can I assume you mean Lisa also signed off that it's trivial to fix the verification code, etc?

Yes. It will require additional test cases to cover the additional functionality. Lisar did not think that was a big issue and she also felt the hardware is now solid to the point where she would be comfortable covering this without having to rely on it being available in terp.

Tim

From: Curtis Abbott [abbott@microunity.com]
Sent: Friday, April 07, 1995 11:32 AM
To: Tim B. Robinson
Subject: something I forgot to mention

Tim B. Robinson wrote (on Fri Apr 7):

Curtis Abbott wrote (on Thu Apr 6):

FYI. Relative to your call on Craig's latest proposal... There's some sentiment beyond Lisa over here that changing the Euterpe spec now would indicate a lack of discipline on our part. I see Craig's proposal changing verification code and perhaps a little bit of our code as well as logic design and layout. I think the impact on the software side is small; my concern is more with the engineering process issues.

Well, this is something he actually raised with a few people including me a good time back but did not publish till we had looked at what it would take. It's actually trivial to implement, so it was hard to argue against including it. I agree that we don't demonstrate as much discipline as we should, but realistically, I don't see this having any impact on tapeout, given we still have bug fixes which must be included.

When you say it's trivial to implement, can I assume you mean Lisa also signed off that it's trivial to fix the verification code, etc?

From: tbr
Sent: Friday, April 07, 1995 2:37 AM
To: Curtis Abbott
Subject: something I forgot to mention

Curtis Abbott wrote (on Thu Apr 6):

FYI. Relative to your call on Craig's latest proposal... There's some sentiment beyond Lisa over here that changing the Euterpe spec now would indicate a lack of discipline on our part. I see Craig's proposal changing verification code and perhaps a little bit of our code as well as logic design and layout. I think the impact on the software side is small; my concern is more with the engineering process issues.

Well, this is something he actually raised with a few people including me a good time back but did not publish till we had looked at what it would take. It's actually trivial to implement, so it was hard to argue against including it. I agree that we don't demonstrate as much discipline as we should, but realistically, I don't see this having any impact on tapeout, given we still have bug fixes which must be included.

Tim

From: pmayer (Patricia Mayer)
Sent: Thursday, April 06, 1995 11:16 PM
To: tbr; pmayer
Cc: dane; yves; rich; arya; graham; rmm; woody; albers; tbe; hestia
Subject: Re: Hestia Main PCB update

> From pmayer Thu Apr 6 17:43:36 1995
> To: tbr
> Subject: Hestia Main PCB update
> Cc: dane, yves, rich, arya, graham, rmm, woody, albers, tbe, hestia
> Content-Length: 1000
>
> Tim,
>
> I've reduced the DRC's on Hestia Main board to 195!
>
> The unconnected pin pairs are 1136. This includes all power and ground
> connections!
>
> There are still 87 unplaced new components.
> Placing the components are dependent on the new mechanical definition
> which is still not available.
>
> *Dane's section has been verified.
> *Jean-Yves Audio and Video sections look good except for the phone
> jack connector. This seems to have swapped connections.
> *I hope to work with Noel tomorrow morning on the power section.
> *Rich will check in with me tomorrow, waiting ECO.
> *Arya will be ready Monday morning for edits. The Balun part is in
> question. If the manufacturer changes, the pinouts will change. The
> RF section is a major area edit.
>
> I don't see meeting the schedule by the end of next week. This will
> probably slip at least a week.
>
> -Pattie
>
>

Did I forget to mention almost all of Euterpe is clean except for the new circuit off the
PLL pins. Thanks to Jay and his support through the PCB crunch.

Sorry Jay and Thanks!
-Pattie

From: Tom Eich [tbe@microunity.com]
Sent: Thursday, April 06, 1995 7:00 PM
To: dbulfer; yves; tbr; pmayer; howard; albers
Cc: pandora
Subject: Pandora is going to be "left-handed"

After reviewing the latest thinking wrt Calliope module with Jean-Yves, I realized that the layout I've created for Pandora needs to change in one simple way: that is that the backplane needs to be on the right side of the box when viewing from the front, with the modules installing through doors on the left, instead of the other way around, as it has been up til now.

This will allow our chips and their heavy heat sinks to remain on the top side of their pcbs, while accomodating Calliope's sensitive I/O routing needs.

The impact of this on the layout work done to date is minimal per Howard. I need to get him a revised criteria for Euterpe pcb to show the change, and then he needs to translate the routing he has done in the x axis to meet the new criteria. I will leave the new mdb file for Dan to input first thing in the morning. Please see me with any questions.

-Tom

Tom Eich
MicroUnity Systems Engineering, Inc.
255 Caspian Dr. Sunnyvale, CA 94089
(408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: Curtis Abbott [abbott@microunity.com]
Sent: Thursday, April 06, 1995 6:10 PM
To: tbr@microunity.com
Subject: followup

Here's a summary of our discussion with some questions. Perhaps you can respond to questions and mistakes. My thought is to send a version of this to you, Scott and Alan as talking points for a meeting or meetings aimed at proposing directions for a cost-effective settop box (ie. mpeg, qam, ntsc).

- Curtis

<intro would go here. need for focus on system level, i.e. both Eu and Ca, and on power & cost. also, desire to leverage design work we've already done.>

Some ideas about Euterpe:

Step 1. go to knob setting 1

- power = 17W
- speed = 250 MHz (?)

Step 2. timing optimize for knob setting 1

- no power change
- speed = 330 MHz (?)

Step 3. change leaf cell family for knob setting 1

- atom count goes down; 3% of total die area (???)
- fewer parasitics so heavily powered nets get faster (what's the bottom line?)

Step 4. change atom for knob setting 1

- remove bias xtors
- other changes?
- atom gets smaller; 3% of total die area (???)

Step 5. SRAM changes for lower speed

- cmos read ports
- cache/buffer loads issue 1/slot; stores 0.5/slot
- reduces power (12/5 W -> ???)
- better software IPC

Step 6. remove 1 Hermes channel

- regain 4% of total die area?

Step 7. rework NB/DRAM interfaces for knob setting 1

- should reduce area, give better (relative) latency

Step 8. add specialized HW and/or new instructions

- e.g. huffman decoder, averaging, rounding, sine wave generation
- goal would be to reduce required instructions for settop benchmark by factor of 2, while still fitting in 1 cm²

From: wingard (Drew Wingard)
Sent: Thursday, April 06, 1995 3:57 PM
To: gmo; microlunatics
Subject: Re: P6 Talk on SITN

Guillermo A. Loyola wrote:

> Did anybody tape the P6 talk given yesterday at the CSL Colloquium,
> and was broadcast this morning at 8am?

Unfortunately, I didn't receive the notice about the P6 talk until late this morning.

Speaking of the Computer Systems Colloquium, here's the one for next week, which will be broadcast next Thursday at 8am:

EE380 Computer Systems Colloquium

Spring Quarter 1994-1995

Lecture #2

Date: Wednesday, Apr 12, 1994

Time: 4:15-5:30 pm

Location: Terman Auditorium

SITN: Thursday, Channel E3, 8:00-9:15

Speaker: Steve McGeady, Intel

Title: The Information Ho Chi Minh Trail: The New Computer
and Communication Industry

ABSTRACT

"We have been bombarded lately with breathless reports about an 'Information Superhighway,' coming soon to your home. But experts can't seem to agree on what devices will connect your home or office to the Infobahn, what services will be offered on it, or what business and technical form this information revolution will take. Intel is at the heart of an industry that has delivered a 1000-fold increase in computing power over the last 15 years. Intel's Communications Technology Lab is now charting a course to integrate the personal computer with the rapidly-increasing communication bandwidth that is being made available to enable new applications in the office and home. Mr. McGeady will outline the history of the shift in the computer industry and analyze the ongoing and impending shifts in the communications industry."

BIOGRAPHY

Steven McGeady is Vice President and General Manager of Intel's Communications Technology Lab, a part of the Intel Architecture Labs. CTL develops advanced software technology and applications that enable personal computers to transmit, receive, and display new types of digital information such as graphics, audio, interactive video, and to participate on high-performance full-service digital networks.

Mr. McGeady has been at Intel for 10 years, leading his group in the creation of the Indeo(TM) Video Software compression technology, the DCI Display Control Interface, key ProShare(TM) Personal Conferencing technology and the Desktop News LAN video capability.

Mr. McGeady studied physics and philosophy at Reed College in Portland, Oregon. His

technical background comes from being an early hacker of
UNIX systems, compilers, operating systems and graphics software.

```
*****
* EE380 is the Computer Systems Laboratory Colloquium.  The Colloquium *
* meets most Wednesdays throughout the academic year.                *
* *                                                                      *
* LECTURES ARE OPEN TO EVERYONE -- FACULTY, STUDENT (ENROLLED OR NOT) *
* INDUSTRIAL VISITORS, OR OTHER INTERESTED PARTIES.                    *
* *                                                                      *
* We frequently have a "dutch treat" dinner for the speaker following *
* the lecture.  If you would like to join us, please contact one of    *
* the course organizers.                                                *
* *                                                                      *
* For information on the class send e-mail with a subject line          *
* mentioning "info" in the subject line to ee380@shasta.stanford.edu.  *
*****
```

From: pmayer (Patricia Mayer)
Sent: Wednesday, April 05, 1995 12:54 PM
To: tbe; dbulfer; woody; howard; philip; tbr; pmayer; lisar
Cc: albers; pandora
Subject: Euterpe PCB Status

This is an updated status of the Pandora Euterpe PCB. Items marked DONE will be removed on the next publication.

Does anyone want to see new updated plots? These can easily be created and distributed. Please move to close your open items. As tbr said "...we need to push this to completion so we can press ahead with the next board..."

> From pmayer Mon Apr 3 12:22:40 1995
> To: tbe, dbulfer, woody, howard, philip, tbr, pmayer, lisar
> Subject: Notes Euterpe PCB Review
> Cc: albers

>
> Netlist edits:
Done > * Woody - New circuit for PLL pins.
Done > * Woody -Add 5v to the Power Connector.
Done > * Woody -Add ground net to connector tabs (flange).
Done > * Woody -Pins offset on half the 168 pin connector.
Done > * Philip - Replace thru hole capacitor for surface mount, 68uf, need part
> number

>
> Verification:
Done > * Howard - The vias for diferential pairs have been moved. (.3 max length
> difference between center lines to outside lines. Not outside to
outside.)

Questionable results- distance between clkIn0 to din0_n7 is .431
dout0_n7 to clkout0 is .250
clkIn1 to din1_n7 is .391
dout1_n7 to clkout1 is .233

> * Lisa - Do we need/want test points? If so are they defined on the
schematic which is desirable for field testing only?
Done > * Has logo been approved? YES! (sort-of)
> * David - dielectric thickness... Please pass info on to Howard for
the Fab
> drawing
Done > * Howard - Verify vias and straight lines meet diferential pairing
Done > * Howard - Maintain 6 mil distance routing even around vias.
New item * Lisa - Logical verification complete befor Fab release

>
> Mechanical:
> * From tbe@microunity.com Thu Mar 30 16:08:52 1995
> Just to clarify, I have only supplied outline criteria and that having to
> do with the Euterpe and SDRAM areas. There are still features such as
> slots and hole to be placed outside the trace and component areas, and I am
> working on completing that design for early next week, but the layout we
> review Monday won't yet have those features in it. Does anyone have a
> problem with reviewing what we will have done by Monday (traces and
> components and pcb outline)?
> * Will have by the end of the week and will incorporate IPC-D-275 requirments
>like component to metal spacing.

>
> Layout issues:
Done > * Howard - Verify vias and straight line routing meet diferential pairing
Done > * Howard - Move capacitors near connector closer to Euterpe.
Done > * Howard - Rotate capacitors around Euterpe so Power pins are closest to
> Euterpe.
Done > * Add ground plane around corner of tab.

Done > * Clear ring of solder mask.
 > * Rename components - Back annotate.
 > * Silkscreen
 > * Drawing dimensions, detail for cutout
 Philip > * Round up special notes and Manufacturing stackup for drawings, Fab
 and Assy
 > * Create Gerbers, drill, ncrouter
 >
 >
 > Plan
 Done > * Howard will continue incorporating edits from review meeting
 Done > * Jay to make netlist changes. Monday
 Done > * Dan to Package. Tuesday
 > If something gets tied up, Howard will move on to Mnemo.
 (Mnemo also dependant on Jay and Dan.) Thursday
 > * Tom to submit Mechanical criteria. (Dan to translate?) Friday
 > * Final edits and photoplots - Tuesday New Item *Final review - Wednesday

 Fabrication will be held until we have Euterpe. We will shoot photoplots for verification,
 especially internal layers and connections.

From: hopper (Mark Hofmann)
Sent: Wednesday, April 05, 1995 4:55 AM
To: vant
Cc: geert (Geert Rosseel); vo (Tom Vo); lisar (Lisa Robinson); tbr (Tim B. Robinson); vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: Re: euterpe lvs results are in

vant writes:

The euterpe lvs results are in...

NUMBER OF UN-MATCHED SCHEMATICS DEVICES	=	0
NUMBER OF UN-MATCHED LAYOUT DEVICES	=	2
NUMBER OF MATCHED SCHEMATICS DEVICES	=	893259
NUMBER OF MATCHED LAYOUT DEVICES	=	893259

This is good, very good. However, there are 604 discrepancy points all associated with open connections. About a third of them are involving opens in BUFSET_ABM_[0123] and the rest I have not tracked down yet.

This is the best results we've had in a long time...

If there are any newer layouts available which have more subblocks, I'd like to run those next. Or if there are things which need to get fixed on this run, then I'll rerun this again.

The results are in:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs
Enjoy...
Thanks,
Dave

This is looking very promising!

The latest top-level euterpe route looks much better. Perhaps after Kurt's rip-up-re-route step we may have a full chip to LVS.

How long did this run take?

thanks,
-hopper

From: vanthof (vant)
Sent: Wednesday, April 05, 1995 11:51 AM
To: geert (Geert Rosseel); hopper (Mark Hofmann); vo (Tom Vo); lisar (Lisa Robinson); tbr (Tim B. Robinson)
Cc: vanthof (Dave Van't Hof); wampler (Kurt Wampler); tom (Tom Laidig)
Subject: euterpe lvs results are in

The euterpe lvs results are in...

NUMBER OF UN-MATCHED SCHEMATICS	DEVICES	=	0
NUMBER OF UN-MATCHED LAYOUT	DEVICES	=	2
NUMBER OF MATCHED SCHEMATICS	DEVICES	=	893259
NUMBER OF MATCHED LAYOUT	DEVICES	=	893259

This is good, very good. However, there are 604 discrepancy points all associated with open connections. About a third of them are involving opens in BUFSET_ABM_[0123] and the rest I have not tracked down yet.

This is the best results we've had in a long time...

If there are any newer layouts available which have more subblocks, I'd like to run those next. Or if there are things which need to get fixed on this run, then I'll rerun this again.

The results are in:

/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs

Enjoy...

Thanks,
Dave

--
Dave Van't Hof MicroUnity Systems Eng., Inc. 255 Caspian Sunnyvale, CA 94089
vanthof@microunity.com 1 408 734-8100
"Don't blame me! I didn't vote for him"

From: hopper (Mark Hofmann)
Sent: Monday, April 03, 1995 5:36 PM
To: Tim B. Robinson
Subject: Re: pif2pim problem

Tim B. Robinson writes:

No problem. Just annoyed it took me so long to think of checking!

Okay. well, I'm glad it wasn't in the pif2pim code. I did touch that today and even though I ran many tests, I'm always kind of nervous when I change things so close to tapeout.

-hopper

From: Tom Eich [tbe@microunity.com]
Sent: Monday, April 03, 1995 8:57 PM
To: pmayer
Cc: dbulfer; woody; howard; philip; tbr; lisar

>> From tbe@microunity.com Mon Apr 3 17:59:55 1995
>> To: pmayer (Patricia Mayer)
>> Subject: Re: Notes Euterpe PCB Review
>> Cc: dbulfer, woody, howard, philip, tbr, lisar
>>
>> Pattie wrote:
>>
>> >snip<
>> >* Tom to submit Mechanical criteria. (Dan to translate?) Thursday
>> >* Howard taking day off. Friday.
>> >* Final edits an photoplots - Tuesday?
>>
>>
>> At the meeting, when I committed to having final criteria by "the end
>> of the week", I meant Friday, but of course that really means Monday.
>> As Howard is taking Friday off, the impact of this may be nil.
>Howard isn't taking Friday off!!
>

Previously, you wrote:

>snip<
>* Howard taking day off. Friday.
>* Final edits an photoplots - Tuesday?

Did I misunderstand this, or did he change his plans?

-Tom

Tom Eich
MicroUnity Systems Engineering, Inc.
255 Caspian Dr. Sunnyvale, CA 94089
(408)734-8100, (408)734-8136 fax

tbe@microunity.com

From: pmayer (Patricia Mayer)
Sent: Monday, April 03, 1995 8:18 PM
To: pmayer; tbe@microunity.com
Cc: dbulfer; woody; howard; philip; tbr; lisar
Subject: Re: Notes Euterpe PCB Review

> From tbe@microunity.com Mon Apr 3 17:59:55 1995
> To: pmayer (Patricia Mayer)
> Subject: Re: Notes Euterpe PCB Review
> Cc: dbulfer, woody, howard, philip, tbr, lisar

> Pattie wrote:

>
> >snip<
> >* Tom to submit Mechanical criteria. (Dan to translate?) Thursday
> >* Howard taking day off. Friday.
> >* Final edits and photoplots - Tuesday?

>
> At the meeting, when I committed to having final criteria by "the end
> of the week", I meant Friday, but of course that really means Monday.
> As Howard is taking Friday off, the impact of this may be nil.
Howard isn't taking Friday off!!

>
>
>Lisar also had raised the issue of having the logical verification
>complete prior to fab release, and this was not assigned in Pattie's
>minutes.

Oops-thnanks for catching that Tom.

>
>
> Thanks,
>
> -Tom

From: gap [gap@charybdis]
Sent: Monday, April 03, 1995 2:50 PM
To: Geert Rosseel; John Mudge; Tim B. Robinson; steve@charybdis
Subject: Meeting with CSM -Wed 2PM

The sales and technical support people for CSM will be here to meet with us on Wednesday at 2:00PM. The preliminary response is that they can provide us the necessary quantity per wafer start of 24 and we provided them with the 100 wafer quantity with two major tape-outs in July and October with each run a split load of wafers. So, on Wednesday we need to be much more specific regarding EPI (or not), expected yield, etc. They provided us with an estimate of about 83 die per wafer gross and a net yield from that gross of approx. 15 die (they currently get 22 die with two levels of metal). As you might guess, ours will be the largest chip they have ever made.

Grant

From: pmayer (Patricia Mayer)
Sent: Monday, April 03, 1995 2:23 PM
To: tbe; dbulfer; woody; howard; philip; tbr; pmayer; lisar
Cc: albers
Subject: Notes Euterpe PCB Review

>From pmayer Fri Mar 31 14:44:49 1995
To: tbe, dbulfer, woody, howard, philip, tbr
Subject: Plots for Euterpe PCB Review
Cc: albers, pmayer

Copies of the current layout state are being distributed for your review.
The meeting is still scheduled for Monday 10:00 in the Engineering conference Room.
There are still many edits required to complete the board. Here's the list so far:

Netlist edits:

- * Woody - New circuit for PLL pins.
- * Woody -Add 5v to the Power Connector.
- * Woody -Add ground net to connector tabs (flange).
- * Woody -Pins offset on half the 168 pin connector.
- * Philip - Replace thru hole capacitor for surface mount, 68uf, need part number

Verification:

- * Howard - The vias for differential pairs have been moved. (.3 max length difference between center lines to outside lines. Not outside to outside.)
- * Lisa - Do we need/want test points? If so are they defined on the schematic which is desirable for field testing only?
- * Has logo been approved? YES! (sort-of)
- * David - dielectric thickness... Please pass info on to Howard for the Fab drawing
- * Howard - Verify vias and straight lines meet differential pairing
- * Howard - Maintain 6 mil distance routing even around vias.

Mechanical:

- * From tbe@microunity.com Thu Mar 30 16:08:52 1995 Just to clarify, I have only supplied outline criteria and that having to do with the Euterpe and SDRAM areas. There are still features such as slots and hole to be placed outside the trace and component areas, and I am working on completing that design for early next week, but the layout we review Monday won't yet have those features in it. Does anyone have a problem with reviewing what we will have done by Monday (traces and components and pcb outline)?
- * Will have by the end of the week and will incorporate IPC-D-275 requirements like component to metal spacing.

Layout issues:

- * Howard - Verify vias and straight line routing meet differential pairing
- * Maintain 6 mil distance routing even around vias.
- * Howard - Move capacitors near connector closer to Euterpe.
- * Howard - Rotate capacitors around Euterpe so Power pins are closest to Euterpe.
- * Add ground plane around corner of tab.
- * Clear ring of solder mask.
- * Rename components - Back annotate.
- * Silkscreen
- * Drawing dimensions, detail for cutout
- * Round up special notes and Manufacturing stackup for drawings, Fab and Assy
- * Create Gerbers, drill, ncrouter

Plan

- * Howard will continue incorporating edits from review meeting
- * Jay to make netlist changes. Monday

- * Dan to Package. Tuesday
- * Howard will contue to edit. If something gets tied up, Howard will move on to Mnemo. (Mnemo also dependant on Jay and Dan.)
- * Tom to submitt Mechanical criteria. (Dan to translate?) Thursday
- * Howard taking day off. Friday.
- * Final edits an photoplots - Tuesday?

From: tbr (Tim B. Robinson)
Sent: Sunday, April 02, 1995 11:46 PM
To: geert (Geert Rosseel)
Cc: billz; dickson; geert; hopper; mws; vo; wampler; woody
Subject: Latest top-level Euterpe

Geert Rosseel wrote (on Sun Apr 2):

Hi,

Can we meet on Monday at 10:00 a.m. to discuss the latest top-level route :

Can we make it sooner? We have a design review of the Pandora Euterpe module scheduled for 10.

1. uu has a lot of disconnects
2. the fat-wire VDDTS obstructs a lot of targets . It didn't used to. What changed ?
3. es-rg interface has a lot of horizontal disconnects : about 3-4 wires per bit for every other byte
4. hcl congestion area is still there

Billz was going to look at this one. If nothing's been done I could look at it. How is hc0?

5. gt congestion area is still there

My attempts at this one so far resulted in timing violations and growth in the sub block. Woody, can we look at this together, I must have done something silly?

6. NB -> XLU bus has +-10 disconnects. Don't understand this .
7. sr-at-cc has about 2 wires per bit disconnected for upper 16 bits : cc needs work.
8. cerberus has a handful of disconnects.
9. disconnects in xlu control

Geert

From: wampler (Kurt Wampler)
Sent: Sunday, April 02, 1995 10:18 PM
To: billz; dickson; geert; hopper; mws; tbr; vo; woody
Subject: Re: Latest top-level Euterpe

Geert writes:

> Can we meet on Monday at 10:00 a.m. to discuss the latest top-level route :

We have a CAD meeting scheduled to talk about seal/crack ring design rule changes at that time. Maybe I can page/swap between the two meetings.

> the fat-wire VDDTS obstructs a lot of targets . It didn't used to.
> What changed ?

Looking at the prior route, VDDTS appears to have been routed as a thin wire by the maze router. Perhaps it failed to route during the fat hwc phase and was finished up by the maze router later, or perhaps it was ripped by the rip-up router and re-routed as a thin wire. (I will protect all hwc nets from being ripped the next time I run the ripper.)

I've examined the target changes we made to avoid the "comb" obstructions; those appear to be correctly modelled, and working as expected.

The pin permutation properties did not appear in the netlist, however, and the geert_euterpe-iter.gplace.nic file does not seem to have been updated to include the increased component flip iteration count and the pin permutation "exitswapsave" command. Let's try to figure out why this didn't work (some subtlety of Makefile dependencies, I suspect).

- Kurt

From: hopper (Mark Hofmann)
Sent: Saturday, April 01, 1995 8:25 AM
To: Tim B. Robinson
Cc: sysadm; tom (Tom Laidig)
Subject: Re: iss lvs failed on euterpe

Tim B. Robinson writes:

Was the old /u/chip/tools version somehow pointing to yet another version?

I can only conclude that was the case, but I don't know. I do know that we had many versions floating around.

-hopper

From: tbr (Tim B. Robinson)
Sent: Saturday, April 01, 1995 12:44 PM
To: geert (Geert Rosseel)
Cc: dickson
Subject: Top-level Impasse ..

Geert Rosseel wrote (on Sat Apr 1):

Hi,

Rich and I believe we have reached some kind of impasse at the top-level. I don't know to resolve this. When Rich runs a block stand-alone, he can get to a result that works. When that block gets included in the top-level power-levels get perturbed and even a relatively minor variation causes rather massive problems at the top-level.

We've tried different things but we don't have a way that really works :

1. When Rich uses the standard top-level Makefile, all the io's that don't go anywhere get pruned.
2. When we use the don't prune option, he gets stuff that should be pruned.

There is an intermediate, which would be to put back the stuff that stopped the top level pruning stuff which was not pruned at the sub-block. As I recall, you took this out because rich had some stuff that was a primary output of the sub-block, but which he intended to have pruned at the top level. I don't know what section that was in, but it may be worth considering putting that stuff back in the makefile and if necessary pruning those gates manually in the srouce.

In fact, given they are sub-block outputs which do not get hooked up at the top level, a way to do this might be just to modify the block interface to eliminate them as outputs. Then the pruning in the sub block will get rid of the gates from the start.

3. We have a way now to control the power-level for the outputs, however power-levels of inputs change depending on what drives these inputs.

We could add something to hold the power of the inputs constant, but I assume that would just either create more timing problems, or simply move the problem to the driving output which would have to change more

Rich did a lot of work on the datapath and even at the top-level, it really looks good, but powerlevels have changed a bit and the edge of es moved by 30 atoms which cause some major overlaps.

Everything is so tight that the slightest variation causes a lot of rework.

The only other radical suggestion I can think of at this point, is to put more of the data path together as a single block. I'm not sure if this would help though. If the problem is that powering up at the top level to meet timing makes things grow, we could be stuck. Is the problem simple growth, or more that things become too ragged?

I don't know how to resolve this, but if we had a better control of the changes between sub-blocks runs and top-level runs, we could finish Euterpe a lot faster.

Seems to me like another major issue here is that we will have a similar problem even if we fix this when we try to iterate the top level.

Has Rich run any of this stuff with the /u/chip/ version of proteus using the or gate pin swapping? One straw to clutch at may be that if he has not, there could be some small improvement from this.

Tim

From: tbr
Sent: Saturday, April 01, 1995 12:24 PM
To: pmayer (Patricia Mayer)
Cc: dbuffer; philip
Subject: Plots for Euterpe PCB Review

Patricia Mayer wrote (on Fri Mar 31):

Copies of the current layout state are being distributed for your review. The meeting is still scheduled for Monday 10:00 in the Engineering conference Room. There are still many edits required to complete the board. Here's the list so far:

The logo looks good to me, and I think we should go with it, and see how it looks when screen printed.

One other thing I may not have mentioned, is while we need to push this to completion so we can press ahead with the next board, we should not release to fab yet. This is because we cannot use the board till we have Euterpe, and there may be problems which need corrections as we get further into the system level design. We need to set a release date by working back from the earliest expected date for Euterpe.

How are the inner planes going to be generated? If you are going to use the composite method as on the Hestia main board (which I think is normal), we should get artworks plotted to verify the flow.

Tim

From: hopper (Mark Hofmann)
Sent: Friday, March 31, 1995 5:06 PM
To: vant
Cc: geert (Geert Rosseel); wampler (Kurt Wampler); vanthof (Dave Van't Hof); lisar (Lisa Robinson); vo (Tom Vo); tbr (Tim B. Robinson)
Subject: Re: euterpe lvs still failed

vant writes:

Hi,

The euterpe lvs came back and there is still a cross in phi_a and phi_b in the gtlb. The cell crclinkt11.ly was last updated on the 27th in the snapshot and the lvs was started on the 28th and also referenced that layout so I'm confident that those edits were picked up. If anyone would like to look at the error file it's in:

`/u/vanthof/compass/mobi/euterpe/tapeout/euterpe.compare/euterpe.lvs`

I'll try to take a look at it in the morning.

Well I'm confused. The edits look good. I don't understand how this layout can give the same result as the last run. Even if we did not identify the error, we surely changed something. Aaargh....

-hopper